HOW DOES SELF-SUPERVISED LEARNING WORK? A REPRESENTATION LEARNING PERSPECTIVE

Anonymous authors
Paper under double-blind review

ABSTRACT

Self-supervised learning (SSL) is a popular machine learning paradigm that utilizes a large amount of unlabeled data to facilitate the learning from a small number of labeled data. While SSL has achieved great success in different tasks, its theoretical understanding remains largely open. In this paper, we aim to theoretically understand a special kind of SSL approaches based on pre-training and fine-tuning. In particular, the SSL approach we consider first trains a neural network based on the unlabeled data with help of pseudo labelers. Then it fine-tunes the pre-trained network on a small amount of labeled data. We prove that, under certain data and neural network models, SSL can achieve nearly zero test loss, while a neural network directly trained by supervised learning on the same amount of labeled data can only achieve constant test loss. Our theoretical result demonstrates a separation between SSL and supervised learning on the same amount of labeled data and sheds light on the essence of representation learning for the success of SSL.

1 INTRODUCTION

With the help of human-annotated labels, supervised learning has achieved remarkable success in several computer vision tasks (Girshick et al., 2014; Long et al., 2015; Krizhevsky et al., 2012; Tran et al., 2015). However, annotating large-scale datasets (e.g., video datasets with temporal dimensions) is time-consuming and costly. In order to reduce the number of labels used for training while maintaining a good prediction performance, a variety of methods have been proposed. Among these methods, *self-supervised learning* (SSL), which leverages unlabeled data to learn feature representations, is one of the most popular and successful approaches. It has been shown to achieve promising prediction accuracy in computer vision (He et al., 2020; Goyal et al., 2019; Chen et al., 2020; Grill et al., 2020; Misra & Maaten, 2020; Caron et al., 2020), speech processing (Baevski et al., 2019; Wu et al., 2020; Baevski & Mohamed, 2020), and natural language processing (NLP) (Devlin et al., 2018; Radford et al., 2018) tasks.

One of the popular self-supervised learning approaches is to utilize a *pseudolabeler* to generate labels for unlabeled data and create "artificial" supervised learning problems known as *pretext tasks*. For example, in image classification, one can first train a teacher network on a small labeled dataset and use it as a *pseudolabeler* to generate pseudo labels for large unlabeled datasets. Then we can train a student network on the combination of labeled and pseudo-labeled images (Xie et al., 2020; Pham et al., 2021b; Rizve et al., 2021). The hope is that by solving these pretext tasks, a good representation can be learned from the unlabeled data that is useful for *downstream learning tasks*. Unfortunately, despite the remarkable success of self-training for neural networks, there has not been a satisfactory theoretical understanding of such approaches.

Some attempts have been made by several recent works to theoretically understand self-supervised learning with *pseudolabeler*. Oymak & Gulcu (2021) considered learning a linear classifier in the Gaussian mixture model setting. They are able to show that in the high dimensional limit, the predictors found by self-training are correlated with the Bayes-optimal predictor. Frei et al. (2022b) further proved that the self-training algorithm can provably converge to the Bayes-optimal predictor for Mixture Models. However, their analyses are limited to the linear classifier, and cannot explain the success of self-supervised learning methods for neural networks.

In this paper, we aim to theoretically explain the success of self-supervised learning from a feature learning perspective. Specifically, we focus on a data model that contains both signal patches and noise patches, where the signal patch is correlated to the label while the noise patch is not. We then consider learning a two-layer convolutional neural network (CNN) with self-supervised learning and provide a comprehensive analysis of the learning process in both pretext and downstream tasks.

The contributions of our work are summarized as follows.

- We theoretically show that the training in the pretext task converges to a global minimum of the training loss. Moreover, we demonstrate that the learned feature representation is highly correlated with the true labels of the data, even though the true labels are unknown and not used during the training on the pretext task.
- Based on our analysis of the learning process for the pretext task, we further show that when fine-tuning the pre-trained model in the downstream task, the final classifier can achieve near-zero test loss and test error. Notably, these guarantees of small test loss and error only require a constant number of labeled training data.
- As a comparison, we show that standard supervised learning cannot learn a good classifier under the same setting. Specifically, we show that, even when the training process converges to a global minimum of the training loss, the learned two-layer CNN can only achieve constant level test loss. This, together with the aforementioned results for self-supervised learning, demonstrates the advantage of self-supervised learning over standard supervised learning.

Notation. We use lower case letters, lower case bold face letters, and upper case bold face letters to denote scalars, vectors, and matrices respectively. For a scalar x, we use $[x]_+$ to denote $\max\{x, 0\}$. For a vector $\mathbf{v} = (v_1, \cdots, v_d)^\top$, we denote by $\|\mathbf{v}\|_2 := \big(\sum_{i=1}^d v_i^2\big)^{\frac{1}{2}}$ its ℓ_2 norm, and use $\text{supp}(\mathbf{v}) := \{j : v_j \neq 0\}$ to denote its support. For two sequences $\{a_k\}$ and $\{b_k\}$, we denote $a_k = O(b_k)$ if $|a_k| \leq C|b_k|$ for some absolute constant C, denote $a_k = \Omega(b_k)$ if $b_k = O(a_k)$, and denote $a_k = \Theta(b_k)$ if $|a_k| \leq C|b_k|$ and $a_k = \Omega(b_k)$. We also denote $a_k = o(b_k)$ if $\lim |a_k/b_k| = 0$. Finally, we use $\widetilde{\Theta}(\cdot)$, $\widetilde{O}(\cdot)$ and $\widetilde{\Omega}(\cdot)$ to omit logarithmic terms in the notations.

2 Related Work

Self-supervised learning (SSL) methods in practice. Many self-supervised learning methods have been proposed to learn visual features from large-scale unlabeled images or videos without using any human annotations. They use various pretext tasks including (1) generation-based ones such as colorizing grayscale images (Zhang et al., 2016), image inpainting (Pathak et al., 2016), image and video generation with GAN (Goodfellow et al., 2014; Brock et al., 2018; Karras et al., 2019; Vondrick et al., 2016; Tulyakov et al., 2018), and (2) context-based ones such as image jigsaw puzzle (Noroozi & Favaro, 2016), geometric transformation (Gidaris et al., 2018; Jing et al., 2018), frame order verification and recognition (Lee et al., 2017; Misra et al., 2016; Wei et al., 2018). Besides applications in computer vision, there are many other natural language processing (NLP) applications of self-supervised learning, such as BERT (Devlin et al., 2018) and its variants (Su et al., 2019; Yang et al., 2019; Kong et al., 2019), as well as GPT (Radford et al., 2018), all of which create auxiliary tasks for next word predictions.

Theory of SSL. Recently, there is a line of theoretical work (Saunshi et al., 2019; Tsai et al., 2020; Mitrovic et al., 2020; Tian et al., 2020b; Wang & Isola, 2020; Tosh et al., 2021b;a) towards understanding *contrastive learning*, which is another mostly used SSL approach based on data augmentation. *Contrastive learning* aims at grouping semantically similar examples, usually generated by augmentations (Chen et al., 2020; Tian et al., 2020a), together, and pushing dissimilar examples far from each other, by optimizing the so-called *contrastive loss* (Hadsell et al., 2006; Tian et al., 2020b). Saunshi et al. (2019) first presented a theoretical framework for analyzing contrastive learning algorithms by introducing latent classes, which characterize semantically similar objects. Tosh et al. (2021b;a) provided theoretical analyses of contrastive learning in both document classification with topic modeling setting and multi-view setting. They showed that contrastive learning is capable of recovering a representation, which reveals useful information to downstream linear models and are nearly optimal. HaoChen et al. (2021) analyzed contrastive learning by using a novel concept called *augmentation graph* on data and proposed a new contrastive loss, which leads to fea-

tures with provable guarantees under linear probe evaluation. Wen & Li (2021) analyzed the feature learning process of contrastive learning for neural network and the effect of data augmentations, while Saunshi et al. (2022) demonstrated that such analyses ignore inductive biases of the function class and training algorithm, cannot adequately explain the success of contrastive learning. Unlike contrastive learning, the theoretical understanding of pretext task-based SSL is still rather limited. The only notable works are Lee et al. (2020) and Wei et al. (2020). Lee et al. (2020) proved generalization guarantees for self-supervised algorithms using empirical risk minimizer of pretext task under certain conditional independence assumption. Wei et al. (2020) proved that under "expansion" assumption, the minimizers of population loss based on self-training and input-consistency regularization will achieve high prediction accuracy.

Pseudo-Labeling. The idea of pseudo-labeling is first given by Lee et al. (2013), where the authors proposed to first train the model on a batch of labeled data, then use the trained model to pseudo-label the unlabeled data, and finally use both labeled data and pseudo-labeled data to train a new model with better performance. Recently, Xie et al. (2020); Pham et al. (2021a) utilized pseudo-labeling to solve semi-supervised computer vision tasks. Specifically, they first trained a teacher model based on the labeled ImageNet data, then used this teacher model to pseudo label the JFT-300M dataset, and finally used the JFT-300M dataset with pseudo labels to train a student model.

Feature learning by neural networks. Our work is also closely related to several recent works that study how neural networks learn the features. Allen-Zhu & Li (2020a) showed that adversarial training purifies the learned features by removing certain "dense mixtures" in the hidden layer weights of the network. Allen-Zhu & Li (2020b) studied how ensemble and knowledge distillation work in deep learning when the data have "multi-view" features. Zou et al. (2021) studied an aspect of feature learning by Adam and GD, and showed that GD can learn the sparse features while Adam may fail even with proper regularization. Notably, there are two concurrent works studying the benign overfitting phenomenon in learning neural networks: Frei et al. (2022a) established theoretical guarantees for benign overfitting of two-layer fully connected neural networks with zero training error and test error close to the Bayes-optimal error, while Cao et al. (2022) studied the benign overfitting phenomenon in training a two-layer convolutional neural network (CNN), achieving arbitrarily small training and test loss. Our work studies a different aspect of feature learning afforded by self-supervised learning versus supervised learning: given a small amount of labeled data, self-supervised learning can learn the features with the help of pretext task (i.e., pseudo-labeler), while supervised learning fails to learn the features and tends to overfit the training data.

3 Problem Setup and Preliminaries

In this section, we first give a brief overview of the pretext-task-based self-supervised learning pipeline. Then we will introduce our data model, the convolutional neural network, and the detail of the training algorithms considered in this paper.

3.1 Pretext-task Based Self-supervised Learning Pipeline

We consider *pretext-based self-supervised learning*, which creates human-designed supervised learning problems to leverage natural structures in the unlabeled data and then utilizes this learned model to improve the final performance.

The typical pipeline of *pretext-based SSL* is shown in Figure 1, which consists of two steps. The first step is *self-supervised pre-training*, in which a predefined *pretext task* is designed, and the intrinsic features of unlabeled data are learned by several feature representation functions. Data $\mathbf{x}_i$ used in pretext tasks are equipped with *pseudo label* $\widehat{y}_i$ generated based on some attributes of the data.

After the training on the pretext task, models that capture features of unlabeled data are obtained, which can be further transferred to *linear downstream tasks* as pre-trained models to improve performance.

3.2 Data distribution

Inspired by recent work (Allen-Zhu & Li, 2020b; Zou et al., 2021; Shen et al., 2022; Cao et al., 2022), we consider a data model where each data input $\mathbf{x}$ consists of two patches $\mathbf{x}^{(1)}$ and $\mathbf{x}^{(2)}$,

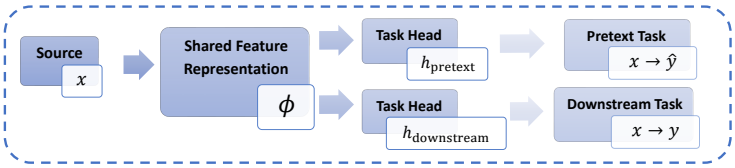

Figure 1: The general pipeline of pretext-based self-supervised learning.

where each patch has d dimensions. We focus on binary classification task, and present our data distribution $\mathcal{D}$ in the following definition.

Definition 3.1. Each data point $(\mathbf{x}, y)$ with $\mathbf{x} = [\mathbf{x}^{(1)\top}, \mathbf{x}^{(2)\top}]^\top \in \mathbb{R}^{2d}$ and $y \in \{-1, +1\}$ is generated as follows: the label y is generated as a Rademacher random variable; one of $\mathbf{x}^{(1)}, \mathbf{x}^{(2)}$ is given by the feature vector $y \cdot \mathbf{v}$, the other is given by a noise vector $\boldsymbol{\xi}$ that is generated from a d-dimensional Gaussian distribution $\mathcal{N}\big(\mathbf{0}, \sigma_p^2(\mathbf{I} - \mathbf{v}\mathbf{v}^\top/\|\mathbf{v}\|_2^2)\big)$. We denote by $\mathcal{D}$ the joint distribution of $(\mathbf{x}, y)$, and denote by $\mathcal{D}_{\mathbf{x}}$ the marginal distribution of $\mathbf{x}$.

The most natural way to think of our data model is to treat patches $\mathbf{x}^{(1)}$ and $\mathbf{x}^{(2)}$ as the embedding of the image data: one of them is a signal which is label-dependent, and the other one is the noise that is label-independent. For simplicity, we assume that the noise patch is generated from the Gaussian distribution $N(\mathbf{0}, \sigma_p^2 \cdot (\mathbf{I} - \mathbf{v}\mathbf{v}^\top \cdot \|\mathbf{v}\|_2^{-2}))$ to ensure that the noise vector is orthogonal to the signal vector $\mathbf{v}$, and only consider the case where the data consists of one signal patch and one noise patch. However, our results and proof techniques can be easily extended to cover the setting with non-orthogonal signal/noise and multiple signal/noise patches. With this simple data model, we aim to show the effectiveness of self-supervised learning and explain the mechanism behind self-supervised learning with neural networks.

Since the positions of signal and noise are not specified in Definition 3.1. It is natural to use a classifier with a convolutional structure that applies the same function to each patch. More specifically, we consider learning a CNN with n_{l} labeled examples $S' = \{(\mathbf{x}_i', y_i')\}_{i=1}^{n_{\mathrm{l}}}$ generated from the distribution $\mathcal{D}$ and n_{u} unlabeled examples $S = \{\mathbf{x}_i\}_{i=1}^{n_{\mathrm{u}}}$ generated from the marginal distribution $\mathcal{D}_{\mathbf{x}}$, where n_l is significantly smaller than the dimension d. If we only use the labeled data, the CNN can easily overfit the training dataset by memorizing the noise patches $\boldsymbol{\xi}_i$. Consequently, the CNN will perform badly on the fresh test data. Therefore, our case is hard to learn without using unlabeled examples.

3.3 Supervised Learning Models

For supervised learning, we consider a two-layer CNN whose filters are applied to the patches $\mathbf{x}^{(1)}$ and $\mathbf{x}^{(2)}$ respectively and parameters in the second layers are set to be ± 1. Then the CNN can be written as $f_{\mathbf{W}}(\mathbf{x}) = f_{\mathbf{W}}^{+1}(\mathbf{x}) - f_{\mathbf{W}}^{-1}(\mathbf{x})$ where $f_{\mathbf{W}}(\mathbf{x})^{+1}, f_{\mathbf{W}}(\mathbf{x})^{-1}$ are formulated as

$$
\begin{aligned}
f_{\mathbf{W}}^{+1}(\mathbf{x}) &= \sum_{j=1}^{m} \Big[\sigma\big(\langle \mathbf{w}_j, \mathbf{x}^{(1)}\rangle\big) + \sigma\big(\langle \mathbf{w}_j, \mathbf{x}^{(2)}\rangle\big)\Big], \\
f_{\mathbf{W}}^{-1}(\mathbf{x}) &= \sum_{j=m+1}^{2m} \Big[\sigma\big(\langle \mathbf{w}_j, \mathbf{x}^{(1)}\rangle\big) + \sigma\big(\langle \mathbf{w}_j, \mathbf{x}^{(2)}\rangle\big)\Big].
\end{aligned}
\tag{3.1}
$$

Here σ is truncated polynomial activation function $\mathrm{ReLU}^3(\cdot) = [\,\cdot\,]_+^3$, m is the width of the network, $\mathbf{w}_j \in \mathbb{R}^d$ denotes the j-th filter, and $\mathbf{W}$ is the collection of all filters $\{\mathbf{w}_j\}_{j=1}^{2m}$. Given labeled training dataset $S' = \{(\mathbf{x}_i', y_i')\}_{i=1}^{n_{\mathrm{l}}}$, we train the CNN model by minimizing the empirical cross-entropy loss

$$
L_{S'}(\mathbf{W}) = \frac{1}{n_{\mathrm{l}}} \sum_{i=1}^{n_{\mathrm{l}}} L_i(\mathbf{W}),
$$

where $L_i(\mathbf{W}) = \ell\big(y_i' \cdot f_{\mathbf{W}}(\mathbf{x}_i')\big)$ with $\ell(z) = \log(1 + \exp(-z))$ denotes the individual loss for the training example $(\mathbf{x}_i, y_i)$. We minimize the empirical function $L_{S'}(\mathbf{W})$ with gradient descent as

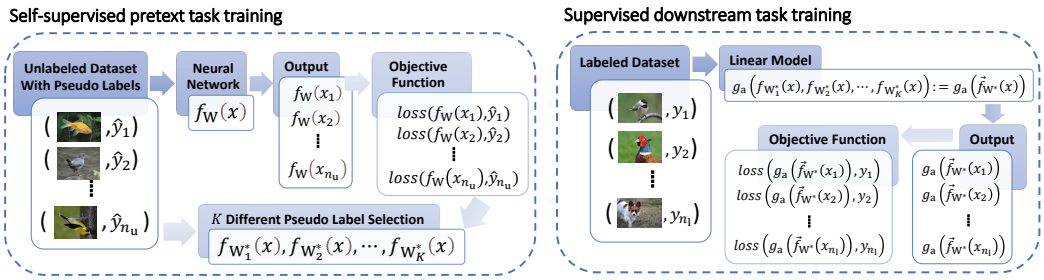

Figure 2: Illustration of our model. The left figure characterizes self-supervised pre-train schema: NN is trained by minimizing errors between pseudo labels $\widehat{y}$ and predictions $f_{\mathbf{W}}(\mathbf{x})$. After self-supervised pretext task training finished, the learned parameters $\{\mathbf{W}_k^*\}_{k=1}^K$ serve as pre-trained models and are transferred to downstream model $g_{\mathbf{a}}$ for finetuning, as shown in the right figure.

follows

$$\mathbf{w}_j^{(t+1)} = \mathbf{w}_j^{(t)} - \eta \cdot \nabla_{\mathbf{w}_j} L_{S'}(\mathbf{W}^{(t)}), \quad \mathbf{w}_j^{(0)} \sim \mathcal{N}(\mathbf{0}, \sigma_0^2 \mathbf{I}), \quad j \in [2m],$$

where $\eta > 0$ is the learning rate and σ_0 defines the scale of random initialization.

3.4 Self-supervised Learning Models

For self-supervised pre-training, we assume that we have access to K pseudolabelers $\{f_k^{\mathrm{w}}\}_{k=1}^K$. The accuracy of k-th pseudolabeler is $p_k \in (1/2, 1)$. Then we use K pseudolabelers to generate K pseudo-labeled dataset $\{S_k\}_{k=1}^K$, where $S_k := \big\{(\mathbf{x}_i, \widehat{y}_{k,i}) \,\big|\, \widehat{y}_{k,i} = f_k^{\mathrm{w}}(\mathbf{x}_i)\big\}_{i=1}^{n_{\mathrm{u}}}$. Next we solve K pretext tasks with two-layer CNN models $\{f_{\mathbf{W}_k}\}_{k=1}^K$ defined in (3.1) using $\{S_k\}_{k=1}^K$ respectively.

We consider learning the model parameter $\mathbf{W}_k$ by optimizing the empirical loss of both pseudo-labeled dataset S_k and labeled dataset $S' = \{(\mathbf{x}_i', y_i')\}_{i=1}^{n_{\mathrm{l}}}$ with weight decay regularization

$$L_{S_k \cup S'}(\mathbf{W}_k) = \frac{1}{n_{\mathrm{u}} + n_{\mathrm{l}}} \left(\sum_{i=1}^{n_{\mathrm{u}}} L_i(\mathbf{W}_k) + \sum_{i'=1}^{n_{\mathrm{l}}} L_{i'}(\mathbf{W}_k) \right) + \frac{\lambda}{2} \|\mathbf{W}_k\|_F^2,$$

where $\lambda \geq 0$ is the regularization parameter, $L_i(\mathbf{W}_k) = \ell\big(\widehat{y}_{k,i} \cdot f_{\mathbf{W}_k}(\mathbf{x}_i)\big)$ denotes the individual loss for the pseudo-labeled data $(\mathbf{x}_i, \widehat{y}_{k,i})$, $L_{i'}(\mathbf{W}_k) = \ell(y_i' \cdot f_{\mathbf{W}_k}(\mathbf{x}_i'))$ denotes the individual loss for the labeled data $(\mathbf{x}_i', y_i')$.

We use gradient descent to minimize the regularized loss function $L_{S_k \cup S'}(\mathbf{W}_k)$. Starting from initial $\mathbf{W}_k^{(0)} := \{\mathbf{w}_{k,j}^{(0)}, j \in [2m]\}$, gradient descent update rule is as follows

$$\mathbf{w}_{k,j}^{(t+1)} = \mathbf{w}_{k,j}^{(t)} - \eta \cdot \nabla_{\mathbf{w}_{k,j}} L_{S_k \cup S'}(\mathbf{W}_k^{(t)}), \quad \mathbf{w}_{k,j}^{(0)} \sim \mathcal{N}(\mathbf{0}, \sigma_0^2 \mathbf{I}_d), \quad j \in [2m], k \in [K]$$

where $\eta > 0$ is the learning rate and σ_0 defines the scale of random initialization.

• **Downstream Task: Linear Model.** The training of pretext task gives us K CNN models with parameters $\{\mathbf{W}_k^*\}_{k=1}^K$. Based on them, for the downstream task, we consider a linear model

$$g_{\mathbf{a}}(\mathbf{x}) = \sum_{k=1}^K a_k f_{\mathbf{W}_k^*}(\mathbf{x}),$$

where $a_k \in \mathbb{R}$ denotes the trainable weight for the k-th pretext model. Then, given $\{f_{\mathbf{W}_k^*}\}_{k=1}^K$ and labeled training data $S' = \{(\mathbf{x}_i', y_i')\}_{i=1}^n$. We consider learning the downstream linear model parameter $\mathbf{a}$ by optimizing the following empirical loss

$$L_{S'}(\mathbf{a}) = \frac{1}{n} \sum_{i=1}^n \ell\big(y_i' \cdot g_{\mathbf{a}}(\mathbf{x}_i')\big).$$

We initialize $\mathbf{a}$ as an all zero vector and optimize the empirical loss by gradient descent with learning rate η, i.e.,

$$\mathbf{a}^{(t+1)} = \mathbf{a}^{(t)} - \eta \cdot \nabla_{\mathbf{a}} L_{S'}(\mathbf{a}^{(t)}), \ \mathbf{a}^{(0)} = \mathbf{0}.$$

4 MAIN RESULTS

In this section, we present the main theoretical results in this paper. We start with a condition that is required by our analysis.

Condition 4.1. We assume the strength of the signal $\|\mathbf{v}\|_2^2 = \Theta(d)$, the noise variance $\sigma_p = \Theta(d^{0.01})$, and the width of the network $m = \text{polylog}(d)$. We also assume that the size of the unlabeled dataset $n_{\text{u}} = \Omega(d^{0.04})$, and labeled data $n_{\text{l}} = \widetilde{\Theta}(1)$. For both supervise learning and self-supervised learning settings, we initialize the weight with $\sigma_0 = \Theta(d^{-3/4})$. For self-supervised learning, we require $\lambda = o(d^{3/4})$ and assume that there exists a constant C such that for all pseudolabelers, their test accuracy $p_k > 1/2 + C$.

Since we generate the noise patch from the Gaussian distribution, the strength of the noise patch is $\|\boldsymbol{\xi}\|_2^2 \approx d^{1.01}$ by standard concentration inequalities, which is larger than the strength of the signal patch $\|\mathbf{v}\|_2^2 = \Theta(d)$. Therefore, Condition 4.1 defines a setting with large noises. The condition of $d \gg n_u \gg n_l$ further ensures that learning is in a sufficiently over-parameterized setting. Here we only require the neural network width m to be polylogarithmic in the dimension d and require the psudolablers to perform better than a random guess.

Theorem 4.2 (Self-supervised Learning: Pretext)**.** Let $k \in [K]$ and consider the self-supervised pre-training of $f_{\mathbf{W}_k}(\mathbf{x})$. For any test data point $(\mathbf{x}, y)$, denote $\widehat{y} = f_k^{\text{w}}(\mathbf{x})$. Then under Condition 4.1, after $T_0 = \widetilde{\Theta}(d^{-\frac{3}{4}}\eta^{-1})$ training iterations with learning rate $\eta = O(d^{-1.1})$, the trained neural network can achieve nearly 0 test error on the distribution $\mathcal{D}$: $\mathbb{P}_{(\mathbf{x},y)\sim\mathcal{D}}[y \cdot f_{\mathbf{W}_k^{(T_0)}}(\mathbf{x}) \leq 0] = o(1)$.

Theorem 4.2 characterizes the prediction power of the feature representation learned in the pretext task using unlabeled data. For any test data point $(\mathbf{x}, y)$, the sign of y can be predicted based on $f_{\mathbf{W}^{(T_0)}}(\mathbf{x})$ with high probability.

Theorem 4.3 (Self-supervised Learning: Downstream)**.** Let $\big\{f_{\mathbf{W}_k^{(T_0^k)}}\big\}_{k=1}^d$ be the neural networks trained according to the d pretext tasks, and consider the learning of the downstream task based in $\big\{f_{\mathbf{W}_k^{(T_0^k)}}\big\}_{k=1}^d$. Under Condition 4.1, after $T' = \widetilde{\Theta}(1)$ iterations with learning rate $\eta = \Theta(1)$, with probability $1 - o(1)$, the obtained $\mathbf{a}^{(T')}$ satisfies:

- Training error is 0: $\frac{1}{n}\sum_{i=1}^n \mathbb{1}[y_i \cdot g_{\mathbf{a}^{(T')}}(\mathbf{x}_i) \leq 0] = 0$.
- Test error and loss are nearly 0: $\mathbb{P}_{(\mathbf{x},y)\sim\mathcal{D}}[y \cdot g_{\mathbf{a}^{(T')}}(\mathbf{x}) \leq 0] = o(1), L_{\mathcal{D}}\big(\mathbf{a}^{(T')}\big) = o(1)$.

Theorem 4.3 shows that the feature representation learned based on the unsupervised pretext tasks can ensure small training and test errors for the supervised downstream task. Notably, this result holds even though we assume that there are only a constant number of labeled data. This shows that self-supervised learning can significantly reduce the need for a large labeled training dataset. For comparison, we also have the following guarantees on the performance of standard supervised learning of CNNs.

Theorem 4.4 (Supervised Learning)**.** Under supervised learning setting, after gradient descent for $T = \widetilde{\Theta}(d^{-0.78}\eta^{-1})$ iterations with learning rate $\eta = O(d^{-1.1})$, then there exists $t \leq T$ such that with probability $1 - o(1)$ the CNNs defined in (3.1) with parameter $\mathbf{W}^{(t)}$ satisfies:

- Training loss is nearly zero: $L_{S'}\big(\mathbf{W}^{(t)}\big) = o(1)$.
- Test loss is high: $L_{\mathcal{D}}\big(\mathbf{W}^{(t)}\big) = \Theta(1)$.

Theorem 4.4 shows that although standard supervised learning can train a CNN model with nearly zero training loss, the obtained CNN model generalizes poorly to test data. Comparing Theorem 4.4 with Theorem 4.3 shows that the generalization of self-supervised learning and supervised learning are largely different. The reason behind this difference is that, the pretext task, with a relatively large number of unlabeled training data, helps learn a feature representation that captures the feature $\mathbf{v}$ in our data model, while direct application of supervised learning can only memorize the noises $\boldsymbol{\xi}_i'$, $i \in [n_l]$ in the training dataset, which is independent of the labels of the data.

A recent line of work Oymak & Gulcu (2021); Frei et al. (2022b) studies the pretext-based self-supervised learning methods with pseudolabelers. Our results are different from theirs in several aspects: (i) we are considering learning with CNNs rather than a linear model, so the problem is highly non-convex with various local minima, which makes the optimization analysis more challenging; (ii) the Bayesian optimal predictor is no longer unique for CNNs. Therefore, we measure the quality of the learned features via downstream task instead of making a comparison with the Bayesian optimal predictor; (iii) They can only deal with the case where the teacher network (pseudolabeler) is the same as the student network (Frei et al., 2022b) or the case where the teacher network (pseudolabeler) is at least as complex as the student network (Oymak & Gulcu, 2021). However, our teacher network (pseudolabeler) is not specified and can be any structure, such as a linear network. Therefore we can handle the case where the student network is more complex than the teacher network, one of the most natural settings for self-training with pseudolabeler (Xie et al., 2020).

5 Proof Sketch

In this section, we present the proof sketch for the self-supervised learning setting. And the proof sketch of the supervised learning setting is given in the appendix.

Pretext task. We consider learning K functions $f_{\mathbf{W}_k}(\mathbf{x})$, $k \in [K]$ based on the pretext task. Since the learning process of these K functions can be analyzed in exactly the same way, here we only focus on the learning of one of these functions. For simplicity of notation, we drop the subscript k in the following proof sketch.

Our study of the pretext task focuses on two aspects of the training process: *feature learning* and *noise memorization*. Specifically, we aim to monitor how the filters in the CNN model learn the feature vector $\mathbf{v}$ and the noise vectors $\boldsymbol{\xi}_i$'s. Therefore, we introduce the following notations.

$$\begin{aligned}
&\widehat{\Lambda}_1^{(t)} := \max_{1\le j\le m} \langle \mathbf{w}_j^{(t)}, \mathbf{v}\rangle,\ \bar{\Lambda}_1^{(t)} := \max_{1\le j\le m} -\langle \mathbf{w}_j^{(t)}, \mathbf{v}\rangle,\\
&\widehat{\Lambda}_{-1}^{(t)} := \max_{m+1\le j\le 2m} -\langle \mathbf{w}_j^{(t)}, \mathbf{v}\rangle,\ \bar{\Lambda}_{-1}^{(t)} := \max_{m+1\le j\le 2m} \langle \mathbf{w}_j^{(t)}, \mathbf{v}\rangle,\\
&\Gamma_i^{(t)} := \max_{1\le j\le 2m} \langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_i\rangle,\ \Gamma_i'^{(t)} := \max_{1\le j\le 2m} \langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_i'\rangle,\ \Gamma^{(t)} = \max\Big\{ \max_{i\in[n_{\mathrm{u}}]} \Gamma_i^{(t)}, \max_{i\in[n_{\mathrm{l}}]} \Gamma_i'^{(t)}, \Big\}.
\end{aligned} \tag{5.1}$$

Based on the above definitions for $r \in \{\pm 1\}$, a larger $\widehat{\Lambda}_r^{(t)}$ implies better feature learning along the positive feature direction $\mathbf{v}$, while a larger $\bar{\Lambda}_r^{(t)}$ implies better feature learning along the negative feature direction $-\mathbf{v}$. Moreover, a larger $\Gamma^{(t)}$ implies a higher level of noise memorization.

Based on the update rule of gradient descent, for the inner products $\langle \mathbf{w}_j^{(t)}, \mathbf{v}\rangle$ and $\langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_l\rangle$, for $j \in [2m]$, $l \in [n_{\mathrm{u}}]$, we can obtain iterative equations in (A.1). With the help of the iterative equations and definitions in (5.1), we can further show the following lemma.

Lemma 5.1. Assume we use both unlabeled data with pseudo labels generated by the pseudolabeler and labeled data for the training of our CNN model. Then for $r \in \{\pm 1\}$, let T_r be the first iteration that $\widehat{\Lambda}_r^{(t)}$ reaches $\Theta(1/m)$, then for $t \in [0, T_r]$, we have

$$\begin{aligned}
&\widehat{\Lambda}_r^{(t+1)} \ge (1-\eta\lambda)\cdot \widehat{\Lambda}_r^{(t)} + \eta\cdot C\cdot \Theta(d)\cdot (\widehat{\Lambda}_r^{(t)})^2, r\in\{\pm 1\},\\
&\bar{\Lambda}_r^{(t+1)} \le (1-\eta\lambda)\cdot \bar{\Lambda}_r^{(t)}, r\in\{\pm 1\},\\
&\Gamma^{(t+1)} \le (1-\eta\lambda)\cdot \Gamma^{(t)} + \eta\cdot \widetilde{\Theta}(d^{0.98})\cdot (\Gamma^{(t)})^2,
\end{aligned}$$

where C is defined in Condition 4.1.

Lemma 5.2. Assume we use only labeled data for the training of our CNN model. Then for $i \in [n_{\mathrm{l}}]$, let T_i' be the first iteration that $\Gamma_i'^{(t)}$ reaches $\Theta(1/m)$, then we have

$$\begin{aligned}
&\widehat{\Lambda}_r^{(t+1)} \le (1-\eta\lambda)\cdot \widehat{\Lambda}_r^{(t)} + \eta\cdot \Theta(d)\cdot \big((\widehat{\Lambda}_r^{(t)})^2 + (\bar{\Lambda}_r^{(t)})^2\big), r\in\{\pm 1\},\\
&\bar{\Lambda}_r^{(t+1)} \le (1-\eta\lambda)\cdot \bar{\Lambda}_r^{(t)}, r\in\{\pm 1\},\\
&\Gamma_i'^{(t+1)} \ge (1-\eta\lambda)\cdot \Gamma_i'^{(t)} + \eta\cdot \widetilde{\Theta}(d^{1.02})\cdot (\Gamma_i'^{(t)})^2, i\in[n_{\mathrm{l}}],\ \text{for } t\in[0, T_i'].
\end{aligned}$$

Based on the results in Lemma 5.1, we can observe that if both pseudo-labeled and labeled data are used for training, the CNN will learn the positive direction of the feature vector $\mathbf{v}$, while barely tending to fit the negative direction of the feature vector or memorize the noise. And if only labeled data are used, the CNN will fit noise faster than a feature, which can be seen from Lemma 5.2. Leveraging Lemmas 5.1 and 5.2, we can obtain following Lemmas 5.3 and 5.4, which characterize the magnitude of feature learning and noise memorization.

Lemma 5.3. If both pseudo-labeled and labeled data are used to train CNN, for $r \in \{\pm 1\}$, let T_r be the first iteration that $\widehat{\Lambda}_r^{(t)}$ reaches $\Theta(1/m)$ respectively. Let $T_0 = \max_{r\in\{\pm 1\}}\{T_r\}$. Then, it holds that $\widehat{\Lambda}_r^{(T_0)} = \widetilde{\Theta}(1)$, $\bar{\Lambda}_r^{(t)} = \widetilde{O}(d^{-\frac{1}{4}})$ and $\Gamma^{(t)} = \widetilde{O}(d^{-0.74})$ for all $t \in [0, T_0]$.

Lemma 5.4. If only labeled data are used to train CNN, for $i \in [n_{\mathrm{l}}]$, let T_i' be the first iteration that $\Gamma_i'^{(t)}$ reaches $\Theta(1/m)$. Let $T_0' = \max_{i\in[n_{\mathrm{l}}]} T_i'$. Then, it holds that $\widehat{\Lambda}_r = \widetilde{O}(d^{-\frac{1}{4}})$, $\bar{\Lambda}_r = \widetilde{O}(d^{-\frac{1}{4}})$ for $r \in \{\pm 1\}$ and $\Gamma_i'^{(t)} = \widetilde{\Theta}(1)$ for $i \in [n_{\mathrm{l}}]$.

The above results indicate the deviation between the two settings. The reason is that assume we consider a sequence $\{x_t\}$ with iterative equation $x_{t+1} = x_t + \eta \cdot C_t x_t^2$. If we only use labeled data, as shown in Lemma 5.2, $\Gamma_i'^{(t)}$ has $C_t = \widetilde{\Theta}(d^{1.02})$ while $\widehat{\Lambda}_r^{(t)}$ has $C_t = \Theta(d)$, therefore $\Gamma_i'^{(t)}$ increases faster than $\widehat{\Lambda}_r^{(t)}$. In contrast, if we use both labeled data and pseudo-labeled data, C_t will be $\widetilde{\Theta}(d^{0.98})$ for $\Gamma_i'^{(t)}$ and $\Theta(d)$ for $\widehat{\Lambda}_r^{(t)}$, leading to a slower increasing speed of $\Gamma_i'^{(t)}$.

Downstream task. After the training on the pretext tasks, we have obtained K classifiers $\{f_{\mathbf{W}_k^{(T_0^k)}}\}_{k=1}^K$. Now we train the second-layer parameters $\mathbf{a}$ with the training data whose true labels are available. The following lemma shows that the l_1-norm of $\mathbf{a}$ will increase with a logarithmic order.

Lemma 5.5. For any learning rate $\eta = \Theta(1)$, we have $\|\mathbf{a}^{(t)}\|_1 = \Theta(\log(t))$. For any labeled data $(\mathbf{x}_i', y_i') \in S'$, we have with high probability that $y_i' \cdot f_{\mathbf{W}^{(t)}}(\mathbf{x}_i') = \|\mathbf{a}^{(t)}\|_1 \cdot \widetilde{\Theta}(1)$. For any newly generated data $(\mathbf{x}, y) \sim \mathcal{D}$, we also have with high probability that $y \cdot f_{\mathbf{W}^{(t)}}(\mathbf{x}) = \|\mathbf{a}^{(t)}\|_1 \cdot \widetilde{\Theta}(1)$.

With the help of the above lemma and note that training error and test error are related to $y \cdot f_{\mathbf{W}^{(T_0)}}(\mathbf{x})$ and test loss is related to $\|\mathbf{a}\|_1$, we can prove that after $T' = \widetilde{\Theta}(1)$ iterations with learning rate $\eta = \Theta(1)$, the model can achieve nearly zero training error, test error, training loss and test loss.

6 EXPERIMENTS

In this section, we perform numerical experiments on synthetic datasets, generated according to Definition 3.1, to verify our main theoretical results. In particular, we set the problem dimension $d = 10000$, labeled training sample size $n_{\mathrm{l}} = 20$ (10 positive samples and 10 negative samples), pseudo-labeled training sample size $n_{\mathrm{u}} = 20000$ (10000 positive samples and 10000 negative samples), feature vector $\mathbf{v}$ sampled from distribution $\mathcal{N}(\mathbf{0}, \boldsymbol{I})$ and noise vector sampled from distribution $\mathcal{N}(\mathbf{0}, \sigma_p^2 \boldsymbol{I})$ where $\sigma_p = 10d^{0.01}$.

For self-supervised learning task, we have a linear pseudolabeler with test error 0.196 ± 0.044. Then, we use this classifier to generate pseudo labels for $n_u = 20000$ unlabeled samples in order to help self-supervised learning. After that, for pretext task, we use these pseudo-labeled samples and n_l labeled samples together to train a CNN with network width $m = 20$, activation function $\sigma(z) = [z]_+^3$, regularization parameter $\lambda = 0.1$ and learning rate $\eta = 1 \times 10^{-4}$. Besides, we initialize CNN parameters from $\mathcal{N}(0, \sigma_0^2)$, where $\sigma_0 = 0.1 \times d^{-3/4}$. After 200 iterations, we can obtain a CNN model with training error close to the error of pseudolabeler and zero test error, according to Table 6. And for a downstream task, we use n_{l} labeled samples to finetune a linear model. By applying learning rate $\eta = 0.1$ and after $T = 100$ iterations, we can obtain a final model with low training and test loss as well as 100% training accuracy and test accuracy.

For supervised learning task, we directly use n_{l} labeled data to train a CNN with network width $m = 20$, activation function $\sigma(z) = [z]_+^3$, learning rate $\eta = 1 \times 10^{-4}$. After 200 iterations, we obtain a CNN with 0 training error and small training loss, about 0.5 test error and high test loss,

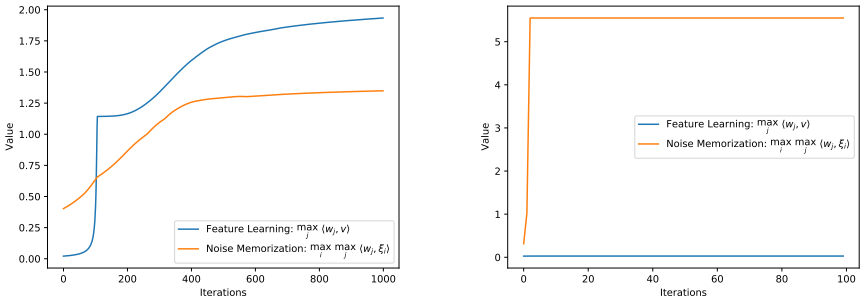

Figure 3: Visualization of the feature learning and noise memorization in the training process. (Left: Self-supervised, Right: Supervised)

	Self-supervised		Supervised
	Pretext	Downstream	
Training error	0.1753±0.0259	0	0
Test error	0	0	0.4982± 0.0208
Training loss	0.4155±0.0418	0.0150±0.0022	$(6.473\pm5.031)\times10^{-7}$
Test loss	0.2200±0.0886	0.0182±0.0021	0.6931±0.0005

Table 1: Training error and loss, test error and loss for self-supervised and supervised learning

which indicates supervised learning will give a model that behaves badly and even no better than a random guess.

Moreover, for synthetic data experiments, we also calculate the inner products $\max_{j\in[m]}\langle \mathbf{w}_j^{(t)}, \mathbf{v}\rangle$ and $\max_{j\in[2m]}\big\{\max_{i\in[n_{\mathrm{u}}]}\langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_i\rangle, \max_{i\in[n_{\mathrm{l}}]}\langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_i'\rangle\big\}$, i.e. $\widehat{\Lambda}_1^{(t)}$ and $\Gamma^{(t)}$, representing feature learning and noise memorization respectively, to verify our key lemmas. The results are reported in Figure 3. It can be seen from Figure 3 that under self-supervised learning setting the algorithm will the feature learning will dominate the noise memorization though the noise patch has a larger norm than the signal patch, while under the supervised learning setting, the algorithm will entirely forget the feature but fit noise. This verifies Lemmas 5.3 and 5.4.

7 Conclusion

In this paper, we provide a theoretical understanding towards the success of self-supervised learning as well as the advantage of self-supervised learning over supervised learning. By considering a simple data model and two-layer CNN, we present a comprehensive analysis of the training procedure from a beyond-NTK feature learning perspective. We prove that the final classifier of a self-supervised learning scenario can achieve near-zero test loss and error with only a constant number of labeled training data, while its supervised-learned counterpart fails to achieve the same performance with the same data complexity.

References

Zeyuan Allen-Zhu and Yuanzhi Li. Feature purification: How adversarial training performs robust deep learning. *arXiv preprint arXiv:2005.10190*, 2020a.

Zeyuan Allen-Zhu and Yuanzhi Li. Towards understanding ensemble, knowledge distillation and self-distillation in deep learning. *arXiv preprint arXiv:2012.09816*, 2020b.

Alexei Baevski and Abdelrahman Mohamed. Effectiveness of self-supervised pre-training for asr. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 7694–7698. IEEE, 2020.

Alexei Baevski, Steffen Schneider, and Michael Auli. vq-wav2vec: Self-supervised learning of discrete speech representations. *arXiv preprint arXiv:1910.05453*, 2019.

Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018.

Yuan Cao, Zixiang Chen, Mikhail Belkin, and Quanquan Gu. Benign overfitting in two-layer convolutional neural networks. *arXiv preprint arXiv:2202.06526*, 2022.

Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in Neural Information Processing Systems*, 33:9912–9924, 2020.

Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607. PMLR, 2020.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

Spencer Frei, Niladri S Chatterji, and Peter L Bartlett. Benign overfitting without linearity: Neural network classifiers trained by gradient descent for noisy linear data. *arXiv preprint arXiv:2202.05928*, 2022a.

Spencer Frei, Difan Zou, Zixiang Chen, and Quanquan Gu. Self-training converts weak learners to strong learners in mixture models. In *International Conference on Artificial Intelligence and Statistics*, pp. 8003–8021. PMLR, 2022b.

Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728*, 2018.

Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 580–587, 2014.

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.

Priya Goyal, Dhruv Mahajan, Abhinav Gupta, and Ishan Misra. Scaling and benchmarking self-supervised visual representation learning. In *Proceedings of the ieee/cvf International Conference on computer vision*, pp. 6391–6400, 2019.

Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent-a new approach to self-supervised learning. *Advances in Neural Information Processing Systems*, 33:21271–21284, 2020.

Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pp. 1735–1742. IEEE, 2006.

Jeff Z HaoChen, Colin Wei, Adrien Gaidon, and Tengyu Ma. Provable guarantees for self-supervised deep learning with spectral contrastive loss. *Advances in Neural Information Processing Systems*, 34, 2021.

Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 9729–9738, 2020.

Longlong Jing, Xiaodong Yang, Jingen Liu, and Yingli Tian. Self-supervised spatiotemporal feature learning via video rotation prediction. *arXiv preprint arXiv:1811.11387*, 2018.

Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4401–4410, 2019.

Lingpeng Kong, Cyprien de Masson d'Autume, Wang Ling, Lei Yu, Zihang Dai, and Dani Yogatama. A mutual information maximization perspective of language representation learning. *arXiv preprint arXiv:1910.08350*, 2019.

Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012.

Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3, pp. 896, 2013.

Hsin-Ying Lee, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. Unsupervised representation learning by sorting sequences. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 667–676, 2017.

Jason D Lee, Qi Lei, Nikunj Saunshi, and Jiacheng Zhuo. Predicting what you already know helps: Provable self-supervised learning. *arXiv preprint arXiv:2008.01064*, 2020.

Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3431–3440, 2015.

Ishan Misra and Laurens van der Maaten. Self-supervised learning of pretext-invariant representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6707–6717, 2020.

Ishan Misra, C Lawrence Zitnick, and Martial Hebert. Shuffle and learn: unsupervised learning using temporal order verification. In *European Conference on Computer Vision*, pp. 527–544. Springer, 2016.

Jovana Mitrovic, Brian McWilliams, Jacob Walker, Lars Buesing, and Charles Blundell. Representation learning via invariant causal mechanisms. *arXiv preprint arXiv:2010.07922*, 2020.

Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European conference on computer vision*, pp. 69–84. Springer, 2016.

Samet Oymak and Talha Cihad Gulcu. Statistical and algorithmic insights for semi-supervised learning with self-training. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2021.

Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2536–2544, 2016.

Hieu Pham, Zihang Dai, Qizhe Xie, and Quoc V Le. Meta pseudo labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11557–11568, 2021a.

Hieu Pham, Zihang Dai, Qizhe Xie, Minh-Thang Luong, and Quoc V. Le. Meta pseudo labels. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021b.

Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training. `https://www.cs.ubc.ca/~amuham01/LING530/papers/radford2018improving.pdf`, 2018.

Mamshad Nayeem Rizve, Kevin Duarte, Yogesh S Rawat, and Mubarak Shah. In defense of pseudo-labeling: An uncertainty-aware pseudo-label selection framework for semi-supervised learning. In *International Conference on Learning Representations (ICLR)*, 2021.

Nikunj Saunshi, Orestis Plevrakis, Sanjeev Arora, Mikhail Khodak, and Hrishikesh Khandeparkar. A theoretical analysis of contrastive unsupervised representation learning. In *International Conference on Machine Learning*, pp. 5628–5637. PMLR, 2019.

Nikunj Saunshi, Jordan Ash, Surbhi Goel, Dipendra Misra, Cyril Zhang, Sanjeev Arora, Sham Kakade, and Akshay Krishnamurthy. Understanding contrastive learning requires incorporating inductive biases. *arXiv preprint arXiv:2202.14037*, 2022.

Ruoqi Shen, Sebastien Bubeck, and Suriya Gunasekar. Data augmentation as feature manipulation. In *International Conference on Machine Learning*, pp. 19773–19808. PMLR, 2022.

Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. Vl-bert: Pre-training of generic visual-linguistic representations. *arXiv preprint arXiv:1908.08530*, 2019.

Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. What makes for good views for contrastive learning? *Advances in Neural Information Processing Systems*, 33:6827–6839, 2020a.

Yuandong Tian, Lantao Yu, Xinlei Chen, and Surya Ganguli. Understanding self-supervised learning with dual deep networks. *arXiv preprint arXiv:2010.00578*, 2020b.

Christopher Tosh, Akshay Krishnamurthy, and Daniel Hsu. Contrastive estimation reveals topic posterior information to linear models. *Journal of Machine Learning Research*, 22(281):1–31, 2021a.

Christopher Tosh, Akshay Krishnamurthy, and Daniel Hsu. Contrastive estimation reveals topic posterior information to linear models. *Journal of Machine Learning Research*, 22(281):1–31, 2021b.

Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pp. 4489–4497, 2015.

Yao-Hung Hubert Tsai, Yue Wu, Ruslan Salakhutdinov, and Louis-Philippe Morency. Demystifying self-supervised learning: An information-theoretical framework. *arXiv preprint arXiv:2006.05576*, 2020.

Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. Mocogan: Decomposing motion and content for video generation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1526–1535, 2018.

Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Generating videos with scene dynamics. *Advances in neural information processing systems*, 29:613–621, 2016.

Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning*, pp. 9929–9939. PMLR, 2020.

Colin Wei, Kendrick Shen, Yining Chen, and Tengyu Ma. Theoretical analysis of self-training with deep networks on unlabeled data. *arXiv preprint arXiv:2010.03622*, 2020.

Donglai Wei, Joseph J Lim, Andrew Zisserman, and William T Freeman. Learning and using the arrow of time. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8052–8060, 2018.

Zixin Wen and Yuanzhi Li. Toward understanding the feature learning process of self-supervised contrastive learning. In *International Conference on Machine Learning*, pp. 11112–11122. PMLR, 2021.

Anne Wu, Changhan Wang, Juan Pino, and Jiatao Gu. Self-supervised representations improve end-to-end speech translation. *arXiv preprint arXiv:2006.12124*, 2020.

Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. Self-training with noisy student improves imagenet classification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10687–10698, 2020.

Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32, 2019.

Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *European conference on computer vision*, pp. 649–666. Springer, 2016.

Difan Zou, Yuan Cao, Yuanzhi Li, and Quanquan Gu. Understanding the generalization of adam in learning neural networks with proper regularization. *arXiv preprint arXiv:2108.11371*, 2021.

A PROOF FOR SELF-SUPERVISED LEARNING SETTING

We consider learning K functions $f_{\mathbf{W}_k}(\mathbf{x})$, $k \in [K]$ based on the pretext task. Since the learning process of these K functions can be analyzed in exactly the same way, here we only focus on the learning of one of these functions. For simplicity of notation, we drop the subscript k in the following proof for section A.1, A.2, A.3, A.4, A.5, A.6 and A.7.

A.1 GRADIENT CALCULATION

Lemma A.1 (Gradient Calculation)**.** The gradient of loss function $L_S(\mathbf{W})$ with respect to weight parameters $\mathbf{w}_j$ is

$$\nabla_{\mathbf{w}_j} L_{S\cup S'}(\mathbf{W}) = -\frac{3}{n_{\mathrm{l}}+n_{\mathrm{u}}} \cdot \bigg(\sum_{i=1}^{n_{\mathrm{u}}} c_i \widehat{y}_i \big([\langle \mathbf{w}_j, y_i \cdot \mathbf{v} \rangle]_+^2 \cdot y_i \cdot \mathbf{v} + [\langle \mathbf{w}_j, \boldsymbol{\xi}_i \rangle]_+^2 \cdot \boldsymbol{\xi}_i \big) + \sum_{i=1}^{n_{\mathrm{l}}} b_i y_i' \big([\langle \mathbf{w}_j, y_i' \cdot \mathbf{v} \rangle]_+^2 \cdot y_i' \cdot \mathbf{v} + [\langle \mathbf{w}_j, \boldsymbol{\xi}_i' \rangle]_+^2 \cdot \boldsymbol{\xi}_i' \big) \bigg) + \lambda \cdot \mathbf{w}_j,$$

for $1 \le j \le m$; and

$$\nabla_{\mathbf{w}_j} L_{S\cup S'}(\mathbf{W}) = \frac{3}{n_{\mathrm{l}}+n_{\mathrm{u}}} \cdot \bigg(\sum_{i=1}^{n_{\mathrm{u}}} c_i \widehat{y}_i \big([\langle \mathbf{w}_j, y_i \cdot \mathbf{v} \rangle]_+^2 \cdot y_i \cdot \mathbf{v} + [\langle \mathbf{w}_j, \boldsymbol{\xi}_i \rangle]_+^2 \cdot \boldsymbol{\xi}_i \big) + \sum_{i=1}^{n_{\mathrm{l}}} b_i y_i' \big([\langle \mathbf{w}_j, y_i' \cdot \mathbf{v} \rangle]_+^2 \cdot y_i' \cdot \mathbf{v} + [\langle \mathbf{w}_j, \boldsymbol{\xi}_i' \rangle]_+^2 \cdot \boldsymbol{\xi}_i' \big) \bigg) + \lambda \cdot \mathbf{w}_j,$$

for $m+1 \le j \le 2m$, where $-\ell'\big(\widehat{y}_i \cdot f_{\mathbf{W}}(\mathbf{x}_i)\big) = \exp\big[-\widehat{y}_i \cdot f_{\mathbf{W}}(\mathbf{x}_i)\big]/(1+\exp\big[-\widehat{y}_i \cdot f_{\mathbf{W}}(\mathbf{x}_i)\big])$ is denoted as c_i and $-\ell'(y_i' \cdot f_{\mathbf{W}}(\mathbf{x}_i')) = \exp[-y_i' \cdot f_{\mathbf{W}}(\mathbf{x}_i')]/(1+\exp[-y_i' \cdot f_{\mathbf{W}}(\mathbf{x}_i')])$ is denoted as b_i.

Proof of Lemma A.1. When $1 \le j \le m$,

$$\begin{aligned}
\nabla_{\mathbf{w}_j} \ell\big(\widehat{y}_i \cdot f_{\mathbf{W}}(\mathbf{x}_i)\big) &= \ell'\big(\widehat{y}_i \cdot f_{\mathbf{W}}(\mathbf{x}_i)\big) \cdot \widehat{y}_i \cdot \nabla_{\mathbf{w}_j} f_{\mathbf{W}}(\mathbf{x}_i) \\
&= -c_i \cdot \widehat{y}_i \cdot \nabla_{\mathbf{w}_j} f_{\mathbf{W}}(\mathbf{x}_i) \\
&= -c_i \widehat{y}_i \cdot \big(\sigma'\big(\langle \mathbf{w}_j, y_i \cdot \mathbf{v} \rangle\big) \cdot y_i \cdot \mathbf{v} + \sigma'\big(\langle \mathbf{w}_j, \boldsymbol{\xi}_i \rangle\big) \cdot \boldsymbol{\xi}_i \big) \\
&= -3 c_i \widehat{y}_i \big([\langle \mathbf{w}_j, y_i \cdot \mathbf{v} \rangle]_+^2 \cdot y_i \cdot \mathbf{v} + [\langle \mathbf{w}_j, \boldsymbol{\xi}_i \rangle]_+^2 \cdot \boldsymbol{\xi}_i \big)
\end{aligned}$$

$$\begin{aligned}
\nabla_{\mathbf{w}_j} \ell\big(y_i' \cdot f_{\mathbf{W}}(\mathbf{x}_i')\big) &= \ell'\big(y_i' \cdot f_{\mathbf{W}}(\mathbf{x}_i')\big) \cdot y_i' \cdot \nabla_{\mathbf{w}_j} f_{\mathbf{W}}(\mathbf{x}_i') \\
&= -b_i \cdot y_i' \cdot \nabla_{\mathbf{w}_j} f_{\mathbf{W}}(\mathbf{x}_i') \\
&= -b_i y_i' \cdot \big(\sigma'(\langle \mathbf{w}_j, y_i' \cdot \mathbf{v} \rangle) \cdot y_i' \cdot \mathbf{v} + \sigma'(\langle \mathbf{w}_j, \boldsymbol{\xi}_i' \rangle) \cdot \boldsymbol{\xi}_i' \big) \\
&= -3 b_i y_i' \cdot \big([\langle \mathbf{w}_j, y_i' \cdot \mathbf{v} \rangle]_+^2 \cdot y_i' \cdot \mathbf{v} + [\langle \mathbf{w}_j, \boldsymbol{\xi}_i' \rangle]_+^2 \cdot \boldsymbol{\xi}_i' \big)
\end{aligned}$$

and when $m+1 \le j \le 2m$,

$$\begin{aligned}
\nabla_{\mathbf{w}_j} \ell\big(\widehat{y}_i \cdot f_{\mathbf{W}}(\mathbf{x}_i)\big) &= 3 c_i \widehat{y}_i \big([\langle \mathbf{w}_j, y_i \cdot \mathbf{v} \rangle]_+^2 \cdot y_i \cdot \mathbf{v} + [\langle \mathbf{w}_j, \boldsymbol{\xi}_i \rangle]_+^2 \cdot \boldsymbol{\xi}_i \big) \\
\nabla_{\mathbf{w}_j} \ell\big(y_i' \cdot f_{\mathbf{W}}(\mathbf{x}_i')\big) &= 3 b_i y_i' \cdot \big([\langle \mathbf{w}_j, y_i' \cdot \mathbf{v} \rangle]_+^2 \cdot y_i' \cdot \mathbf{v} + [\langle \mathbf{w}_j, \boldsymbol{\xi}_i' \rangle]_+^2 \cdot \boldsymbol{\xi}_i' \big)
\end{aligned}$$

Note that $\nabla_{\mathbf{w}_j} L_{S\cup S'}(\mathbf{W}) = \big(\sum_{i=1}^{n_{\mathrm{u}}} \nabla_{\mathbf{w}_j} \ell\big(\widehat{y}_i \cdot f_{\mathbf{W}}(\mathbf{x}_i)\big) + \sum_{i=1}^{n_{\mathrm{l}}} \nabla_{\mathbf{w}_j} \ell\big(y_i' \cdot f_{\mathbf{W}}(\mathbf{x}_i')\big) \big)/(n_{\mathrm{l}}+n_{\mathrm{u}}) + \lambda \cdot \mathbf{w}_j$, we have proved the lemma. □

A.2 INNER PRODUCT UPDATE RULE CALCULATION

When the model is trained by gradient descent, the update rule can be formulated by

$$\mathbf{w}_j^{(t+1)} = \mathbf{w}_j^{(t)} - \eta \cdot \nabla_{\mathbf{w}_j} L_S(\mathbf{W}^{(t)}), \quad j \in [2m]. \tag{A.1}$$

We study the performance of entire training process from two perspective: feature learning and noise memorization. Mathematically, we will focus on two quantities: $\langle \mathbf{w}_j^{(t)}, \mathbf{v} \rangle$ and $\langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_l \rangle$. And then we have following lemma for the inner product update rule.

Lemma A.2 (Inner Product Update Rule)**.** The feature learning and noise memorization performance of gradient descent can be formulated by

$$\langle \mathbf{w}_j^{(t+1)}, \mathbf{v} \rangle = (1-\eta\lambda) \cdot \langle \mathbf{w}_j^{(t)}, \mathbf{v} \rangle + \frac{3\eta u_j}{n_{\mathrm{l}} + n_{\mathrm{u}}} \bigg(\sum_{i=1}^{n_{\mathrm{u}}} y_i \widehat{y}_i c_i^{(t)} [\langle \mathbf{w}_j^{(t)}, y_i \cdot \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 + \sum_{i=1}^{n_{\mathrm{l}}} b_i^{(t)} [\langle \mathbf{w}_j^{(t)}, y_i' \cdot \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 \bigg),$$

$$\langle \mathbf{w}_j^{(t+1)}, \boldsymbol{\xi}_l \rangle = (1-\eta\lambda) \cdot \langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_l \rangle + \frac{3\eta u_j}{n_{\mathrm{l}} + n_{\mathrm{u}}} \bigg(\sum_{i=1}^{n_{\mathrm{u}}} \widehat{y}_i c_i^{(t)} [\langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_i \rangle]_+^2 \langle \boldsymbol{\xi}_i, \boldsymbol{\xi}_l \rangle + \sum_{i=1}^{n_{\mathrm{l}}} y_i' b_i^{(t)} [\langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_i' \rangle]_+^2 \langle \boldsymbol{\xi}_i', \boldsymbol{\xi}_l \rangle \bigg),$$

$$\langle \mathbf{w}_j^{(t+1)}, \boldsymbol{\xi}_l' \rangle = (1-\eta\lambda) \cdot \langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_l' \rangle + \frac{3\eta u_j}{n_{\mathrm{l}} + n_{\mathrm{u}}} \bigg(\sum_{i=1}^{n_{\mathrm{u}}} \widehat{y}_i c_i^{(t)} [\langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_i \rangle]_+^2 \langle \boldsymbol{\xi}_i, \boldsymbol{\xi}_l' \rangle + \sum_{i=1}^{n_{\mathrm{l}}} y_i' b_i^{(t)} [\langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_i' \rangle]_+^2 \langle \boldsymbol{\xi}_i', \boldsymbol{\xi}_l' \rangle \bigg),$$

where $j \in [2m]$, $l \in [n_{\mathrm{u}}]$ and $u_j := \mathbb{1}_{[1 \le j \le m]} - \mathbb{1}_{[m+1 \le j \le 2m]}$.

Proof of Lemma A.2. According to lemma A.1 and gradient descent update rule A.1, we have

$$\mathbf{w}_j^{(t+1)} = (1-\eta\lambda) \cdot \mathbf{w}_j^{(t)} + \frac{3\eta u_j}{n_{\mathrm{l}} + n_{\mathrm{u}}} \cdot \bigg(\sum_{i=1}^{n_{\mathrm{u}}} c_i \widehat{y}_i \big([\langle \mathbf{w}_j, y_i \cdot \mathbf{v} \rangle]_+^2 \cdot y_i \cdot \mathbf{v} + [\langle \mathbf{w}_j, \boldsymbol{\xi}_i \rangle]_+^2 \cdot \boldsymbol{\xi}_i \big) + \sum_{i=1}^{n_{\mathrm{l}}} b_i y_i' \big([\langle \mathbf{w}_j, y_i' \cdot \mathbf{v} \rangle]_+^2 \cdot y_i' \cdot \mathbf{v} + [\langle \mathbf{w}_j, \boldsymbol{\xi}_i' \rangle]_+^2 \cdot \boldsymbol{\xi}_i' \big) \bigg)$$

Taking inner product with feature vector $\mathbf{v}$ and noise patch $\boldsymbol{\xi}_l$ and note that $\mathbf{v}$ is orthogonal to $\boldsymbol{\xi}_l$ according to the data model, we have

$$\begin{aligned}
\langle \mathbf{w}_j^{(t+1)}, \mathbf{v} \rangle &= (1-\eta\lambda) \cdot \langle \mathbf{w}_j^{(t)}, \mathbf{v} \rangle + \frac{3\eta u_j}{n_{\mathrm{l}} + n_{\mathrm{u}}} \cdot \bigg(\sum_{i=1}^{n_{\mathrm{u}}} c_i^{(t)} \widehat{y}_i \big([\langle \mathbf{w}_j, y_i \cdot \mathbf{v} \rangle]_+^2 y_i \|\mathbf{v}\|_2^2 + [\langle \mathbf{w}_j, \boldsymbol{\xi}_i \rangle]_+^2 \langle \boldsymbol{\xi}_i, \mathbf{v} \rangle \big) \\
&\quad + \sum_{i=1}^{n_{\mathrm{l}}} b_i^{(t)} y_i' \big([\langle \mathbf{w}_j, y_i' \cdot \mathbf{v} \rangle]_+^2 y_i' \|\mathbf{v}\|_2^2 + [\langle \mathbf{w}_j, \boldsymbol{\xi}_i' \rangle]_+^2 \langle \boldsymbol{\xi}_i', \mathbf{v} \rangle \big) \bigg) \\
&= (1-\eta\lambda) \cdot \langle \mathbf{w}_j^{(t)}, \mathbf{v} \rangle + \frac{3\eta u_j}{n_{\mathrm{l}} + n_{\mathrm{u}}} \bigg(\sum_{i=1}^{n_{\mathrm{u}}} y_i \widehat{y}_i c_i^{(t)} [\langle \mathbf{w}_j^{(t)}, y_i \cdot \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 \\
&\quad + \sum_{i=1}^{n_{\mathrm{l}}} b_i^{(t)} [\langle \mathbf{w}_j^{(t)}, y_i' \cdot \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 \bigg),
\end{aligned}$$

$$\begin{aligned}
\langle \mathbf{w}_j^{(t+1)}, \boldsymbol{\xi}_l \rangle &= (1-\eta\lambda) \cdot \langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_l \rangle + \frac{3\eta u_j}{n_{\mathrm{l}} + n_{\mathrm{u}}} \bigg(\sum_{i=1}^{n_{\mathrm{u}}} c_i^{(t)} \widehat{y}_i \big([\langle \mathbf{w}_j, y_i \cdot \mathbf{v} \rangle]_+^2 y_i \langle \mathbf{v}, \boldsymbol{\xi}_l \rangle + [\langle \mathbf{w}_j, \boldsymbol{\xi}_i \rangle]_+^2 \langle \boldsymbol{\xi}_i, \boldsymbol{\xi}_l \rangle \big) \\
&\quad + \sum_{i=1}^{n_{\mathrm{l}}} b_i^{(t)} y_i' \big([\langle \mathbf{w}_j, y_i' \cdot \mathbf{v} \rangle]_+^2 y_i' \langle \mathbf{v}, \boldsymbol{\xi}_l \rangle + [\langle \mathbf{w}_j, \boldsymbol{\xi}_i' \rangle]_+^2 \langle \boldsymbol{\xi}_i', \boldsymbol{\xi}_l \rangle \big) \bigg) \\
&= (1-\eta\lambda) \cdot \langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_l \rangle + \frac{3\eta u_j}{n_{\mathrm{l}} + n_{\mathrm{u}}} \bigg(\sum_{i=1}^{n_{\mathrm{u}}} \widehat{y}_i c_i^{(t)} [\langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_i \rangle]_+^2 \langle \boldsymbol{\xi}_i, \boldsymbol{\xi}_l \rangle \\
&\quad + \sum_{i=1}^{n_{\mathrm{l}}} y_i' b_i^{(t)} [\langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_i' \rangle]_+^2 \langle \boldsymbol{\xi}_i', \boldsymbol{\xi}_l \rangle \bigg),
\end{aligned}$$

and

$$\langle \mathbf{w}_j^{(t+1)}, \boldsymbol{\xi}_l' \rangle = (1-\eta\lambda)\cdot \langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_l' \rangle + \frac{3\eta u_j}{n_{\mathrm{l}}+n_{\mathrm{u}}}\bigg(\sum_{i=1}^{n_{\mathrm{u}}} \widehat{y}_i c_i^{(t)} [\langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_i \rangle]_+^2 \langle \boldsymbol{\xi}_i, \boldsymbol{\xi}_l' \rangle + \sum_{i=1}^{n_{\mathrm{l}}} y_i' b_i^{(t)} [\langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_i' \rangle]_+^2 \langle \boldsymbol{\xi}_i', \boldsymbol{\xi}_l' \rangle \bigg),$$

which completes the proof. □

A.3 Estimate $\widehat{\Lambda}_r^{(0)}, \bar{\Lambda}_r^{(0)}, \Gamma_i^{(0)}, \Gamma_i'^{(0)}$

Let $\widehat{\Lambda}_1^{(t)} = \max_{1\le j\le m} \langle \mathbf{w}_j^{(t)}, \mathbf{v}\rangle$, $\widehat{\Lambda}_{-1}^{(t)} = \max_{m+1\le j\le 2m} -\langle \mathbf{w}_j^{(t)}, \mathbf{v}\rangle$, $\bar{\Lambda}_1^{(t)} = \max_{m+1\le j\le 2m} \langle \mathbf{w}_j^{(t)}, \mathbf{v}\rangle$, $\bar{\Lambda}_{-1}^{(t)} = \max_{1\le j\le m} -\langle \mathbf{w}_j^{(t)}, \mathbf{v}\rangle$, which characterize the *feature learning* aspect of training process. An easy way to distinguish between $\widehat{\Lambda}_r^{(t)}$ and $\bar{\Lambda}_r^{(t)}$ is that $\widehat{\Lambda}_r^{(t)}$ should be large while $\bar{\Lambda}_r^{(t)}$ should be small.

Let $\Gamma_i^{(t)} = \max_{1\le j\le 2m} \langle \mathbf{w}_j, \boldsymbol{\xi}_i \rangle, i \in [n_{\mathrm{u}}]$, $\Gamma_i'^{(t)} = \max_{1\le j\le 2m} \langle \mathbf{w}_j, \boldsymbol{\xi}_i' \rangle, i \in [n_{\mathrm{l}}]$, which characterize the *noise memorization* aspect of training process with respect to a particular sample.

Let $\Gamma^{(t)} = \max\big\{ \max_{i\in[n_{\mathrm{u}}]} \Gamma_i^{(t)}, \max_{i\in[n_{\mathrm{l}}]} \Gamma_i'^{(t)} \big\}$, which characterize the *noise memorization* aspect of training process regardless of the sample index.

We first provide the concentration inequality for $\widehat{\Lambda}_r^{(0)}$ and $\bar{\Lambda}_r^{(0)}$ in the following lemma.

Lemma A.3. With probability at least $1-4\delta$ with respect to the randomness of initialization of $\mathbf{w}$,we have

$$\big|\widehat{\Lambda}_r^{(0)} - \mathbb{E}[\widehat{\Lambda}_r^{(0)}]\big| < \sqrt{8\log\Big(\frac{1}{\delta}\Big)}\sigma_0\|\mathbf{v}\|_2,$$

$$\big|\bar{\Lambda}_r^{(0)} - \mathbb{E}[\bar{\Lambda}_r^{(0)}]\big| < \sqrt{8\log\Big(\frac{1}{\delta}\Big)}\sigma_0\|\mathbf{v}\|_2,$$

and

$$\mathbb{E}[\widehat{\Lambda}_r^{(0)}] \asymp \sqrt{\log(m)}\sigma_0\|\mathbf{v}\|_2, \mathbb{E}[\bar{\Lambda}_r^{(0)}] \asymp \sqrt{\log(m)}\sigma_0\|\mathbf{v}\|_2, r\in\{\pm 1\}.$$

Proof of lemma A.3. Note that $\widehat{\Lambda}_1^{(0)} = \max_{1\le j\le m}\langle \mathbf{w}_j^{(0)}, \mathbf{v}\rangle, \widehat{\Lambda}_{-1}^{(0)} = \max_{m+1\le j\le 2m} -\langle \mathbf{w}_j^{(0)}, \mathbf{v}\rangle$, $\bar{\Lambda}_1^{(0)} = \max_{m+1\le j\le 2m}\langle \mathbf{w}_j^{(0)}, \mathbf{v}\rangle$ and $\bar{\Lambda}_{-1}^{(0)} = \max_{m+1\le j\le 2m} -\langle \mathbf{w}_j^{(0)}, \mathbf{v}\rangle$, $\mathbf{w}_j^{(0)} \sim \mathcal{N}(\mathbf{0}, \sigma_0^2\mathbf{I})$ and $\mathbf{v}$ is a fixed vector. Therefore, $\langle \mathbf{w}_j^{(0)}, \mathbf{v}\rangle \sim \mathcal{N}(0, \sigma_0^2\|\mathbf{v}\|_2^2)$, $-\langle \mathbf{w}_j^{(0)}, \mathbf{v}\rangle \sim \mathcal{N}(0, \sigma_0^2\|\mathbf{v}\|_2^2)$ for all $1\le j\le 2m$ and $\widehat{\Lambda}_r^{(0)}, \bar{\Lambda}_r^{(0)}, r\in\{\pm 1\}$ are identically distributed. Therefore, without loss of generality, we only need to discuss the concentration of $\widehat{\Lambda}_1^{(0)}$. Then, we introduce following theorem, which is commonly used in the study of Gaussian processes, to discuss the concentration of $\widehat{\Lambda}_1^{(0)}, \widehat{\Lambda}_{-1}^{(0)}$.

Theorem A.4 (Borell-TIS inequality)**.** Let X be a centered Gaussian on $\mathbb{R}^m$ and set $\sigma_X^2 := \max_{i\in[m]} E(X_i^2)$. Then for each $t>0$,

$$\mathbb{P}\bigg(\Big| \max_{i\in[m]} X_i - \mathbb{E}\big(\max_{i\in[m]} X_i\big)\Big| > t \bigg) \le 2e^{-\frac{t^2}{2\sigma_X^2}}.$$

By applying this theorem, we have

$$\mathbb{P}\Big(\big|\widehat{\Lambda}_1^{(0)} - \mathbb{E}[\widehat{\Lambda}_1^{(0)}]\big| > t \Big) \le 2e^{-\frac{t^2}{2\sigma_0^2\|\mathbf{v}\|_2^2}}.$$

For the expectation of $\widehat{\Lambda}_r^{(0)}$ and $\bar{\Lambda}_r^{(0)}$, we give the following lemma.

Lemma A.5. Let $Y = \max_{1\le i\le m} X_i$, where $X_i \sim \mathcal{N}(0,\sigma^2)$ are i.i.d. random variables. Then

$$\frac{1}{\sqrt{\pi\log 2}}\sigma\sqrt{\log m} \le \mathbb{E}[Y] \le \sqrt{2}\sigma\sqrt{\log m}.$$

By applying the lemma above, we have

$$\mathbb{E}[\widehat{\Lambda}_1^{(0)}] \asymp \sqrt{\log(m)}\sigma_0\|\mathbf{v}\|_2,$$

which completes the proof.

□

Then we provide concentration inequality for $\Gamma_i^{(0)}$ in the following lemma.

Lemma A.6. Suppose that $d \geq \Omega(\log(m(n_{\mathrm{u}}+n_{\mathrm{l}})/\delta))$, $m = \Omega(\log(1/\delta))$. Then with probability at least $1-\delta$,

$$\frac{\sigma_0\sigma_p\sqrt{d}}{4} \leq \Gamma_i^{(0)} \leq 2\sqrt{\log(16m(n_{\mathrm{u}}+n_{\mathrm{l}})/\delta)}\cdot\sigma_0\sigma_p\sqrt{d}, \text{ for all } i\in[n_{\mathrm{u}}],$$
$$\frac{\sigma_0\sigma_p\sqrt{d}}{4} \leq \Gamma_i'^{(0)} \leq 2\sqrt{\log(16m(n_{\mathrm{u}}+n_{\mathrm{l}})/\delta)}\cdot\sigma_0\sigma_p\sqrt{d}, \text{ for all } i\in[n_{\mathrm{l}}].$$

Proof of Lemma A.6. By Lemma C.1, with probability at least $1-\delta/4$,

$$\begin{aligned}\sigma_p\sqrt{d}/\sqrt{2} \leq \|\boldsymbol{\xi}_i\|_2 \leq \sqrt{3/2}\cdot\sigma_p\sqrt{d}, \text{ for } i\in[n_{\mathrm{u}}],\\ \sigma_p\sqrt{d}/\sqrt{2} \leq \|\boldsymbol{\xi}_i'\|_2 \leq \sqrt{3/2}\cdot\sigma_p\sqrt{d}, \text{ for } i\in[n_{\mathrm{l}}].\end{aligned} \tag{A.2}$$

Therefore, by Gaussian tail bound and union bound, with probability at least $1-\delta/4$,

$$\begin{aligned}\langle\mathbf{w}_j^{(0)},\boldsymbol{\xi}_i\rangle \leq |\langle\mathbf{w}_j^{(0)},\boldsymbol{\xi}_i\rangle| \leq \sqrt{2\log(8m/\delta)}\cdot\sigma_0\|\boldsymbol{\xi}_i\|_2, \text{ for } i\in[n_{\mathrm{u}}],\\ \langle\mathbf{w}_j^{(0)},\boldsymbol{\xi}_i'\rangle \leq |\langle\mathbf{w}_j^{(0)},\boldsymbol{\xi}_i'\rangle| \leq \sqrt{2\log(8m/\delta)}\cdot\sigma_0\|\boldsymbol{\xi}_i'\|_2, \text{ for } i\in[n_{\mathrm{l}}].\end{aligned} \tag{A.3}$$

Note that $\mathbb{P}\big(\sigma_0\sigma_p\sqrt{d}/4 > \langle\mathbf{w}_j^{(0)},\boldsymbol{\xi}_i\rangle\big)$ is an absolute constant and therefore by the condition on m, we have

$$\begin{aligned}\mathbb{P}\bigg(\frac{\sigma_0\sigma_p\sqrt{d}}{4} \leq \Gamma_i^{(t)}\bigg) &= \mathbb{P}\bigg(\frac{\sigma_0\sigma_p\sqrt{d}}{4} \leq \max_{j\in[2m]}\langle\mathbf{w}_j^{(0)},\boldsymbol{\xi}_i\rangle\bigg)\\ &= 1-\mathbb{P}\bigg(\frac{\sigma_0\sigma_p\sqrt{d}}{4} > \max_{j\in[2m]}\langle\mathbf{w}_j^{(0)},\boldsymbol{\xi}_i\rangle\bigg)\\ &= 1-\bigg(\mathbb{P}\bigg(\frac{\sigma_0\sigma_p\sqrt{d}}{4} > \langle\mathbf{w}_j^{(0)},\boldsymbol{\xi}_i\rangle\bigg)\bigg)^{2m}\\ &\geq 1-\frac{\delta}{4},\end{aligned}$$

and

$$\mathbb{P}\bigg(\frac{\sigma_0\sigma_p\sqrt{d}}{4} \leq \Gamma_i'^{(t)}\bigg) \geq 1-\frac{\delta}{4}.$$

On the other hand, according to (A.2) and (A.3), we have

$$\begin{aligned}&\mathbb{P}\big(\Gamma_i^{(t)} \leq 2\sqrt{\log(16m(n_{\mathrm{u}}+n_{\mathrm{l}})/\delta)}\cdot\sigma_0\sigma_p\sqrt{d}\big)\\ &= \mathbb{P}\Big(\max_{j\in[2m]}\langle\mathbf{w}_j^{(0)},\boldsymbol{\xi}_i\rangle \leq 2\sqrt{\log(16m(n_{\mathrm{u}}+n_{\mathrm{l}})/\delta)}\cdot\sigma_0\sigma_p\sqrt{d}\Big)\\ &\geq 1-\frac{\delta}{4},\end{aligned}$$

and

$$\mathbb{P}\big(\Gamma_i'^{(t)} \leq 2\sqrt{\log(16m(n_{\mathrm{u}}+n_{\mathrm{l}})/\delta)}\cdot\sigma_0\sigma_p\sqrt{d}\big) \geq 1-\frac{\delta}{4},$$

which completes the proof. □

A.4 STAGE I OF GD: ON-DIAGONAL FEATURE LEARNING

In this stage, $\widehat{\Lambda}_1^{(t)}$ and $\widehat{\Lambda}_{-1}^{(t)}$ respectively increase to magnitude $\Theta(1/m)$ and $\bar{\Lambda}_1^{(t)}$, $\bar{\Lambda}_{-1}^{(t)}$ and $\Gamma_j^{(t)}$ remain small, the same magnitude as initialization. In order to characterize the behaviour of feature learning and noise memorization during Stage I, we decompose the analysis into following three parts:

1. First, in lemma A.11, we provide a lower bound of the update rules of on-diagonal feature learning term of $\widehat{\Lambda}_1^{(t)}, \widehat{\Lambda}_{-1}^{(t)}$ to lower-bound their increasing speed, and an upper bound of off-diagonal feature learning term $\bar{\Lambda}_1^{(t)}, \bar{\Lambda}_{-1}^{(t)}$ to indicate their decrease.
2. Second, in lemma A.13, we provide a upper bound of the update rules of noise memorization term $\Gamma^{(t)}$ to upper-bound its increasing speed.
3. Third, we provide a useful lemma, which is a derivation of Claim C.20 in Allen-Zhu & Li (2020b), which we call tensor power method. By applying tensor power method, we will prove that:
   • When $\widehat{\Lambda}_1^{(t)}$ reaches $\Theta(1/m)$ at T_1, $\bar{\Lambda}_1^{(t)}$ and $\Gamma^{(t)}$ remain a magnitude no more than initialization.
   • When $\widehat{\Lambda}_{-1}^{(t)}$ reaches $\Theta(1/m)$ at T_{-1}, $\bar{\Lambda}_{-1}$ and $\Gamma^{(t)}$ remain a magnitude no more than initialization.

A.4.1 UPPER BOUND AND LOWER BOUND FOR $\widehat{\Lambda}_1^{(t)}, \widehat{\Lambda}_{-1}^{(t)}$ AND $\bar{\Lambda}_1^{(t)}, \bar{\Lambda}_{-1}^{(t)}$

We first consider part I of Stage I of GD when $\max_{r\in\{+1,-1\}} \big\{\widehat{\Lambda}_r^{(t)}, \bar{\Lambda}_r^{(t)}\big\} \leq \Theta(m^{-1})$.

In this stage, we first prove following lemma:

Lemma A.7. As long as $\max_{r\in\{+1,-1\}} \big\{\widehat{\Lambda}_r^{(t)}, \bar{\Lambda}_r^{(t)}\big\} \leq \Theta(m^{-1})$, we have $c_i^{(t)} := -\ell'\big(\widehat{y}_i \cdot f_{\mathbf{W}^{(t)}}(\mathbf{x}_i)\big)$ and $b_i^{(t)} := -\ell'\big(y_i' \cdot f_{\mathbf{W}^{(t)}}(\mathbf{x}_i')\big)$ remains $1/2 \pm o(1)$.

Proof of lemma A.7. Note that $\ell(z) = \log(1+\exp(-z))$ and $-\ell'(z) = \exp(-z)/\big(1+\exp(-z)\big)$, and without loss of generality assuming $\widehat{y}_i = y_i = 1$, we can express $c_i^{(t)}$ as follow:

$$c_i^{(t)} = -\ell'(f_{\mathbf{W}^{(t)}}(\mathbf{x}_i)) = \frac{e^{\sum_{j=m+1}^{2m}[\sigma(\langle \mathbf{w}_j^{(t)}, \mathbf{v}\rangle)+\sigma(\langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_i\rangle)]}}{e^{\sum_{j=1}^{m}[\sigma(\langle \mathbf{w}_j^{(t)}, \mathbf{v}\rangle)+\sigma(\langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_i\rangle)]} + e^{\sum_{j=m+1}^{2m}[\sigma(\langle \mathbf{w}_j^{(t)}, \mathbf{v}\rangle)+\sigma(\langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_i\rangle)]}},$$

Since $\sigma(\langle \mathbf{w}_j^{(t)}, \mathbf{v}\rangle)$ dominates $\sigma(\langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}\rangle)$ for $j \in [m]$, which will be proved later by using *tensor power method*, we have

$$c_i^{(t)} = \frac{e^{\sum_{j=m+1}^{2m}[\sigma(\langle \mathbf{w}_j^{(t)}, \mathbf{v}\rangle)+\sigma(\langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_i\rangle)]}}{e^{\sum_{j=1}^{m}\sigma(\langle \mathbf{w}_j^{(t)}, \mathbf{v}\rangle)+\{\text{lower order term}\}} + e^{\sum_{j=m+1}^{2m}[\sigma(\langle \mathbf{w}_j^{(t)}, \mathbf{v}\rangle)+\sigma(\langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_i\rangle)]}}.$$

On the one side,

$$c_i^{(t)} \geq \frac{1}{e^{\sum_{j=1}^{m}\sigma(\langle \mathbf{w}_j^{(t)}, \mathbf{v}\rangle)+\{\text{lower order term}\}}+1} \geq \frac{1}{e^{m(\widehat{\Lambda}_1^{(t)})^3}+1} \geq \frac{1}{e^{\Theta(m^{-2})}+1} = \frac{1}{2+o(1)} = \frac{1}{2} - o(1).$$

On the other side, according to lemma 5.3, we have $\bar{\Lambda}_1^{(t)} = \widetilde{O}(d^{-\frac{1}{4}})$ and $\Gamma^{(t)} = \widetilde{O}(d^{-\frac{1}{4}+\epsilon})$, it follows that

$$\begin{aligned} c_i^{(t)} &\leq \frac{e^{m(\bar{\Lambda}_1^{(t)})^3+m(\Gamma^{(t)})^3}}{e^{\sum_{j=1}^{m}\sigma(\langle \mathbf{w}_j^{(t)}, \mathbf{v}\rangle)+\{\text{lower order term}\}} + e^{m(\bar{\Lambda}_1^{(t)})^3+m(\Gamma^{(t)})^3}} \\ &= \frac{1+o(1)}{e^{\sum_{j=1}^{m}\sigma(\langle \mathbf{w}_j^{(t)}, \mathbf{v}\rangle)+\{\text{lower order term}\}} + 1 + o(1)} \\ &\leq \frac{1+o(1)}{1+1+o(1)} = \frac{1}{2} + o(1). \end{aligned}$$

Therefore, we have $c_i^{(t)} = 1/2 \pm o(1)$ if $\widehat{y}_i = y_i = 1$ and other cases ($\widehat{y}_i = y_i = 1$, $\widehat{y}_i = -y_i$, $b_i^{(t)}$) can be proved in a similar way. □

By applying above lemma, we can obtain following lemma:

Lemma A.8. For any $\delta < 1/2$, with probability at least $1 - 2\delta$ over pseudo labels generated by the pseudolabeler, we have

$$\left|\frac{1}{n_{\mathrm{u}}}\sum_{i=1}^{n_{\mathrm{u}}}\widehat{y}_i y_i c_i^{(t)} - \Big(p - \frac{1}{2}\Big)\right| < \sqrt{\frac{1}{8n_{\mathrm{u}}}\log\frac{1}{\delta}} + o(1),$$

where $o(1)$ is with respect to d.

If we denote $\{(\mathbf{x}_i, y_i)|y_i = 1, i \in [n_{\mathrm{u}}]\}$ as S_1, $\{(\mathbf{x}_i, y_i)|y_i = -1, i \in [n_{\mathrm{u}}]\}$ as S_{-1}, $|S_1|$ as n_1 and $|S_{-1}|$ as n_{-1}, we have with probability at least $1 - 4\delta$ that

$$\left|\frac{1}{n_1}\sum_{i=1}^{n_1}\widehat{y}_i y_i c_i^{(t)} - \Big(p - \frac{1}{2}\Big)\right| < \sqrt{\frac{1}{8n_1}\log\frac{1}{\delta}} + o(1),$$

and

$$\left|\frac{1}{n_{-1}}\sum_{i=1}^{n_{-1}}\widehat{y}_i y_i c_i^{(t)} - \Big(p - \frac{1}{2}\Big)\right| < \sqrt{\frac{1}{8n_{-1}}\log\frac{1}{\delta}} + o(1).$$

Proof of lemma A.8. First, according to lemma A.7, we have

$$\frac{1}{n_{\mathrm{u}}}\sum_{i=1}^{n_{\mathrm{u}}}\widehat{y}_i y_i c_i^{(t)} = \frac{1}{n_{\mathrm{u}}}\sum_{i=1}^{n_{\mathrm{u}}}\widehat{y}_i y_i\Big(c_i^{(t)} - \frac{1}{2}\Big) + \frac{1}{2n_{\mathrm{u}}}\sum_{i=1}^{n_{\mathrm{u}}}\widehat{y}_i y_i = \frac{1}{2n_{\mathrm{u}}}\sum_{i=1}^{n_{\mathrm{u}}}\widehat{y}_i y_i \pm o(1) \tag{A.4}$$

Then, according to Hoeffding's inequality when $a_i = -1, b_i = 1$, we have

$$\mathbb{P}\bigg(\bigg|\frac{1}{n_{\mathrm{u}}}\sum_{i=1}^{n_{\mathrm{u}}}\widehat{y}_i y_i - \mathbb{E}\Big[\frac{1}{n_{\mathrm{u}}}\sum_{i=1}^{n_{\mathrm{u}}}\widehat{y}_i y_i\Big]\bigg| \geq t\bigg) \leq 2\exp\bigg(-\frac{2n_{\mathrm{u}}^2 t^2}{\sum_{i=1}^{n_{\mathrm{u}}}(a_i - b_i)^2}\bigg) = 2\exp\left(-2n_{\mathrm{u}}t^2\right).$$

Note that pseudo label $\widehat{y}_i$ generated by pseudolabeler takes y_i with probability p and $-y_i$ with probability $1 - p$, we have $\mathbb{E}\big[\frac{1}{n_{\mathrm{u}}}\sum_{i=1}^{n_{\mathrm{u}}}\widehat{y}_i y_i\big] = \frac{1}{n_{\mathrm{u}}}\sum_{i=1}^{n_{\mathrm{u}}}\mathbb{E}\big[\widehat{y}_i y_i\big] = 2p - 1$. It follows that

$$\mathbb{P}\bigg(\bigg|\frac{1}{2n_{\mathrm{u}}}\sum_{i=1}^{n_{\mathrm{u}}}\widehat{y}_i y_i - \Big(p - \frac{1}{2}\Big)\bigg| \geq t\bigg) \leq 2\exp\left(-8n_{\mathrm{u}}t^2\right),$$

and therefore

$$\left|\frac{1}{2n_{\mathrm{u}}}\sum_{i=1}^{n_{\mathrm{u}}}\widehat{y}_i y_i - \Big(p - \frac{1}{2}\Big)\right| < \sqrt{\frac{1}{8n_{\mathrm{u}}}\log\frac{1}{\delta}} \tag{A.5}$$

holds with probability at least $1 - 2\delta$. According to (A.4) and (A.5), we have

$$\left|\frac{1}{2n_{\mathrm{u}}}\sum_{i=1}^{n_{\mathrm{u}}}\widehat{y}_i y_i - \Big(p - \frac{1}{2}\Big)\right| < \sqrt{\frac{1}{8n_{\mathrm{u}}}\log\frac{1}{\delta}} + o(1),$$

which verifies the first statement of the lemma. And the other part of the lemma can be proved in a similar way. □

According to above lemma and note that $n_{\mathrm{u}}, n_1, n_{-1} = \omega(1)$, we have further that

$$\left|\frac{1}{n_{\mathrm{u}}}\sum_{i=1}^{n_{\mathrm{u}}}\widehat{y}_i y_i c_i^{(t)} - \Big(p - \frac{1}{2}\Big)\right| = o(1),\ \left|\frac{1}{n_r}\sum_{i=1}^{n_r}\widehat{y}_i y_i c_i^{(t)} - \Big(p - \frac{1}{2}\Big)\right| = o(1), r \in \{\pm 1\}, \tag{A.6}$$

with high probability.

Besides, we also need an approximation about n_1 and n_{-1}, which is given as the following lemma:

Lemma A.9. For $r \in \{\pm 1\}$, it holds with probability at least $1 - 2\delta$ that

$$\left|n_r - \frac{n_{\mathrm{u}}}{2}\right| < \sqrt{\frac{n_{\mathrm{u}}}{2}\log\frac{1}{\delta}},$$

where $n_r := |\{(\mathbf{x}_i, y_i)|y_i = r, i \in [n_{\mathrm{u}}]\}|$.

Proof of lemma A.9. Note that $n_r = \sum_{i=1}^{n_\mathrm{u}} \mathbb{1}[X_i = r], r \in \{\pm 1\}$ where X_i takes label $+1$ or -1 with equal probability $1/2$, according to Hoeffding's inequality, we have

$$\mathbb{P}\Big(\Big|\sum_{i=1}^{n_\mathrm{u}} \mathbb{1}[X_i = r] - \mathbb{E}\Big[\sum_{i=1}^{n_\mathrm{u}} \mathbb{1}[X_i = r]\Big]\Big| \geq t\Big) \leq 2\exp\Big(-\frac{2t^2}{n_\mathrm{u}}\Big), r \in \{\pm 1\},$$

and it follows that

$$\mathbb{P}\Big(\Big|n_r - \frac{n_\mathrm{u}}{2}\Big| \geq t\Big) \leq 2\exp\Big(-\frac{2t^2}{n_\mathrm{u}}\Big), r \in \{\pm 1\},$$

leading to

$$\Big|n_r - \frac{n_\mathrm{u}}{2}\Big| < \sqrt{\frac{n_\mathrm{u}}{2}\log\frac{1}{\delta}},$$

with probability at least $1 - 2\delta$. □

For labeled dataset $S' = \{(\mathbf{x}'_i, y'_i)\}_{i=1}^{n_\mathrm{l}}$, we also have

Lemma A.10. For $r \in \{\pm 1\}$, it holds with probability at least $1 - 2\delta$ that

$$\Big|n'_r - \frac{n_\mathrm{l}}{2}\Big| < \sqrt{\frac{n_\mathrm{l}}{2}\log\frac{1}{\delta}},$$

where $n'_r := |\{(\mathbf{x}'_i, y'_i) | y'_i = r, i \in [n_\mathrm{l}]\}|$.

Then we are prepared to estimate an lower bound of increasing speed of $\widehat{\Lambda}^{(t)}$ and an upper bound of decreasing speed of $\bar{\Lambda}^{(t)}$ in the following lemma.

Lemma A.11. For $\widehat{\Lambda}_1^{(t)} := \max_{1\leq j\leq m}\langle \mathbf{w}_j^{(t)}, \mathbf{v}\rangle$ and $\widehat{\Lambda}_{-1}^{(t)} := \max_{m+1\leq j\leq 2m}\langle \mathbf{w}_j^{(t)}, -\mathbf{v}\rangle$, we have with high probability that

$$\widehat{\Lambda}_r^{(t+1)} \geq (1 - \eta\lambda)\cdot\widehat{\Lambda}_r^{(t)} + \eta\cdot\Big(p - \frac{1}{2}\Big)\cdot\Theta(d)\cdot(\widehat{\Lambda}_r^{(t)})^2, r \in \{\pm 1\}.$$

For $\bar{\Lambda}_1^{(t)} := \max_{m+1\leq j\leq 2m}\langle \mathbf{w}_j^{(t)}, \mathbf{v}\rangle$ and $\bar{\Lambda}_1^{(t)} := \max_{1\leq j\leq m}\langle \mathbf{w}_j^{(t)}, -\mathbf{v}\rangle$, we have with high probability that

$$\bar{\Lambda}_r^{(t+1)} \leq (1 - \eta\lambda)\cdot\bar{\Lambda}_r^{(t)}, r \in \{\pm 1\}.$$

Proof of lemma A.11. We first prove the former inequality. Let $j^* = \arg\max_{1\leq j\leq m}\langle \mathbf{w}_j^{(t)}, \mathbf{v}\rangle$ and note that $u_{j^*} = \mathbb{1}_{[1\leq j\leq m]} - \mathbb{1}_{[m+1\leq j\leq 2m]} = 1$, then we have

$$\begin{aligned}
&\widehat{\Lambda}_1^{(t+1)} \geq \langle \mathbf{w}_{j^*}^{(t+1)}, \mathbf{v}\rangle \\
&= (1 - \eta\lambda)\cdot\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v}\rangle + \frac{3\eta}{n_\mathrm{l} + n_\mathrm{u}}\Bigg(\underbrace{\sum_{i=1}^{n_\mathrm{u}} y_i\widehat{y}_i c_i^{(t)}[\langle \mathbf{w}_{j^*}^{(t)}, y_i\cdot\mathbf{v}\rangle]_+^2\|\mathbf{v}\|_2^2}_{\clubsuit} + \underbrace{\sum_{i=1}^{n_\mathrm{l}} b_i^{(t)}[\langle \mathbf{w}_{j^*}^{(t)}, y'_i\cdot\mathbf{v}\rangle]_+^2\|\mathbf{v}\|_2^2}_{\bigstar}\Bigg)
\end{aligned}$$

Then we respectively estimate terms $\clubsuit$ and $\bigstar$.

For $\clubsuit$, note the definition of j^* that $\widehat{\Lambda}_1^{(t)} = \langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v}\rangle$ and note the increasing property of $\widehat{\Lambda}_1^{(t)}$ and $\widehat{\Lambda}_1^{(0)} > 0$ with high probability, we have $\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v}\rangle > 0$. It follows that

$$\begin{aligned}
\underbrace{\sum_{i=1}^{n_\mathrm{u}} y_i\widehat{y}_i c_i^{(t)}[\langle \mathbf{w}_{j^*}^{(t)}, y_i\cdot\mathbf{v}\rangle]_+^2\|\mathbf{v}\|_2^2}_{\clubsuit} &= \sum_{i\in S_1} y_i\widehat{y}_i c_i^{(t)}[\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v}\rangle]_+^2\|\mathbf{v}\|_2^2 + \sum_{i\in S_{-1}} y_i\widehat{y}_i c_i^{(t)}[-\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v}\rangle]_+^2\|\mathbf{v}\|_2^2 \\
&= \sum_{i\in S_1} y_i\widehat{y}_i c_i^{(t)}[\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v}\rangle]_+^2\|\mathbf{v}\|_2^2 \\
&= \Big(\sum_{i\in S_1} y_i\widehat{y}_i c_i^{(t)}\Big)\cdot\|\mathbf{v}\|_2^2\cdot\big(\widehat{\Lambda}_1^{(t)}\big)^2 \\
&= n_1\cdot\Big(p - \frac{1}{2} \pm o(1)\Big)\cdot\|\mathbf{v}\|_2^2\cdot\big(\widehat{\Lambda}_1^{(t)}\big)^2, \qquad \text{(A.7)}
\end{aligned}$$

where $S_1 := \{(\mathbf{x}_i, y_i) | y_i = 1, i \in [n_\mathrm{u}]\}$, $S_{-1} := \{(\mathbf{x}_i, y_i) | y_i = -1, i \in [n_\mathrm{u}]\}$, $n_1 = |S_1|$ and the last equality is due to (A.6).

For $\bigstar$, similarly we have

$$\begin{aligned}
\underbrace{\sum_{i=1}^{n_\mathrm{l}} b_i^{(t)} [\langle \mathbf{w}_{j^*}^{(t)}, y_i' \cdot \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2}_{\bigstar} &= \sum_{i \in S_1'} b_i^{(t)} [\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 + \sum_{i \in S_{-1}'} b_i^{(t)} [-\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 \\
&= \sum_{i \in S_1'} b_i^{(t)} [\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 \\
&= \bigg(\sum_{i \in S_1'} b_i^{(t)} \bigg) \cdot \|\mathbf{v}\|_2^2 \cdot \big(\widehat{\Lambda}_1^{(t)} \big)^2 \\
&= n_1' \cdot \Big(\frac{1}{2} \pm o(1) \Big) \cdot \|\mathbf{v}\|_2^2 \cdot \big(\widehat{\Lambda}_1^{(t)} \big)^2, \qquad \text{(A.8)}
\end{aligned}$$

where $S_1' = \{(\mathbf{x}_i', y_i') | y_i' = 1, i \in [n_\mathrm{l}]\}$, $S_{-1}' = \{(\mathbf{x}_i', y_i') | y_i' = -1, i \in [n_\mathrm{l}]\}$, $n_1' = |S_1'|$ and the last equality is due to lemma A.7.

According to (A.7) and (A.8), we have

$$\begin{aligned}
&\widehat{\Lambda}_1^{(t+1)} \\
&\geq (1 - \eta\lambda) \cdot \widehat{\Lambda}_1^{(t)} + \frac{3\eta}{n_\mathrm{l} + n_\mathrm{u}} \bigg(n_1 \cdot \Big(p - \frac{1}{2} \pm o(1) \Big) \cdot \|\mathbf{v}\|_2^2 \cdot \big(\widehat{\Lambda}_1^{(t)} \big)^2 + n_1' \cdot \Big(\frac{1}{2} \pm o(1) \Big) \cdot \|\mathbf{v}\|_2^2 \cdot \big(\widehat{\Lambda}_1^{(t)} \big)^2 \bigg) \\
&= (1 - \eta\lambda) \cdot \widehat{\Lambda}_1^{(t)} + \frac{3\eta n_1}{n_\mathrm{l} + n_\mathrm{u}} \cdot \Big(p - \frac{1}{2} \pm o(1) \Big) \cdot \|\mathbf{v}\|_2^2 \cdot \big(\widehat{\Lambda}_1^{(t)} \big)^2 + \frac{3\eta n_1'}{n_\mathrm{l} + n_\mathrm{u}} \cdot \Big(\frac{1}{2} \pm o(1) \Big) \cdot \|\mathbf{v}\|_2^2 \cdot \big(\widehat{\Lambda}_1^{(t)} \big)^2 \\
&= (1 - \eta\lambda) \cdot \widehat{\Lambda}_1^{(t)} + 3\eta \cdot \bigg(\frac{n_1}{n_\mathrm{l} + n_\mathrm{u}} \cdot \Big(p - \frac{1}{2} \pm o(1) \Big) + \frac{n_1'}{n_\mathrm{l} + n_\mathrm{u}} \cdot \Big(\frac{1}{2} \pm o(1) \Big) \bigg) \cdot \|\mathbf{v}\|_2^2 \cdot \big(\widehat{\Lambda}_1^{(t)} \big)^2 \\
&= (1 - \eta\lambda) \cdot \widehat{\Lambda}_1^{(t)} + 3\eta \cdot \bigg(\underbrace{\frac{n_1}{n_\mathrm{l} + n_\mathrm{u}} \cdot \Big(p - \frac{1}{2} \Big) + \frac{n_1'}{n_\mathrm{l} + n_\mathrm{u}} \cdot \frac{1}{2}}_{\spadesuit} \pm o(1) \bigg) \cdot \|\mathbf{v}\|_2^2 \cdot \big(\widehat{\Lambda}_1^{(t)} \big)^2. \qquad \text{(A.9)}
\end{aligned}$$

According to lemma A.9 and lemma A.10, and note that $n_\mathrm{l} = \widetilde{\Theta}(1), n_\mathrm{u} = \omega(d^{4\epsilon})$, we have for $\spadesuit$ that with probability at least $1 - 4\delta$

$$\begin{aligned}
&\bigg| \underbrace{\frac{n_1}{n_\mathrm{l} + n_\mathrm{u}} \cdot \Big(p - \frac{1}{2} \Big) + \frac{n_1'}{n_\mathrm{l} + n_\mathrm{u}} \cdot \frac{1}{2}}_{\spadesuit} - \frac{n_\mathrm{u}}{2(n_\mathrm{l} + n_\mathrm{u})} \cdot \Big(p - \frac{1}{2} \Big) - \frac{n_\mathrm{l}}{2(n_\mathrm{l} + n_\mathrm{u})} \cdot \frac{1}{2} \bigg| \\
&\leq \frac{|n_1 - \frac{n_\mathrm{u}}{2}|}{n_\mathrm{l} + n_\mathrm{u}} \cdot \Big(p - \frac{1}{2} \Big) + \frac{|n_1' - \frac{n_\mathrm{l}}{2}|}{n_\mathrm{l} + n_\mathrm{u}} \cdot \frac{1}{2} \\
&\leq \frac{\sqrt{\frac{n_\mathrm{u}}{2} \log \frac{1}{\delta}}}{n_\mathrm{l} + n_\mathrm{u}} \cdot \Big(p - \frac{1}{2} \Big) + \frac{\sqrt{\frac{n_\mathrm{l}}{2} \log \frac{1}{\delta}}}{n_\mathrm{l} + n_\mathrm{u}} \cdot \frac{1}{2} \\
&= \Theta\Big(\frac{1}{\sqrt{n_\mathrm{u}}} \Big) \\
&= o(1)
\end{aligned}$$

Therefore, note that $n_\mathrm{u} = \omega(n_\mathrm{l})$ and $n_\mathrm{u} = \omega(1)$, we have

$$\begin{aligned}
\underbrace{\frac{n_1}{n_\mathrm{l} + n_\mathrm{u}} \cdot \Big(p - \frac{1}{2} \Big) + \frac{n_1'}{n_\mathrm{l} + n_\mathrm{u}} \cdot \frac{1}{2}}_{\spadesuit} &= \frac{n_\mathrm{u}}{2(n_\mathrm{l} + n_\mathrm{u})} \cdot \Big(p - \frac{1}{2} \Big) + \frac{n_\mathrm{l}}{2(n_\mathrm{l} + n_\mathrm{u})} \cdot \frac{1}{2} \pm o(1) \\
&= \frac{1}{2} \cdot \Big(p - \frac{1}{2} \Big) \pm o(1) \qquad \text{(A.10)}
\end{aligned}$$

Plugging (A.10) into (A.9), we have

$$\begin{aligned}\widehat{\Lambda}_1^{(t+1)} &\geq (1-\eta\lambda)\cdot\widehat{\Lambda}_1^{(t)} + 3\eta\cdot\left(\frac{1}{2}\cdot\left(p-\frac{1}{2}\right)\pm o(1)\right)\cdot\|\mathbf{v}\|_2^2\cdot\big(\widehat{\Lambda}_1^{(t)}\big)^2 \\ &= (1-\eta\lambda)\cdot\widehat{\Lambda}_1^{(t)} + \eta\cdot\left(p-\frac{1}{2}\right)\cdot\Theta(d)\cdot\big(\widehat{\Lambda}_1^{(t)}\big)^2, \end{aligned} \tag{A.11}$$

which verifies the first inequality of case $r=1$ in the lemma.

Let $j^{**} = \mathrm{argmax}_{m+1\leq j\leq 2m}\langle \mathbf{w}_j^{(t)}, -\mathbf{v}\rangle$ and note that $u_{j^{**}} = \mathbb{1}_{[1\leq j\leq m]} - \mathbb{1}_{[m+1\leq j\leq 2m]} = -1$, we have

$$\begin{aligned}&\widehat{\Lambda}_{-1}^{(t+1)} \geq \langle \mathbf{w}_{j^*}^{(t+1)}, -\mathbf{v}\rangle \\ &= (1-\eta\lambda)\cdot\langle \mathbf{w}_{j^{**}}^{(t)}, -\mathbf{v}\rangle + \frac{3\eta}{n_{\mathrm{l}}+n_{\mathrm{u}}}\Bigg(\underbrace{\sum_{i=1}^{n_{\mathrm{u}}} y_i\widehat{y}_i c_i^{(t)}[\langle \mathbf{w}_{j^{**}}^{(t)}, y_i\cdot\mathbf{v}\rangle]_+^2\|\mathbf{v}\|_2^2}_{\clubsuit} \\ &\qquad + \underbrace{\sum_{i=1}^{n_{\mathrm{l}}} b_i^{(t)}[\langle \mathbf{w}_{j^{**}}^{(t)}, y_i'\cdot\mathbf{v}\rangle]_+^2\|\mathbf{v}\|_2^2}_{\bigstar}\Bigg)\end{aligned}$$

For $\clubsuit$, note the definition of j^{**} that $\widehat{\Lambda}_{-1}^{(t)} = \langle \mathbf{w}_{j^{**}}^{(t)}, -\mathbf{v}\rangle$ and note the increasing property of $\widehat{\Lambda}_{-1}^{(t)}$ and $\widehat{\Lambda}_{-1}^{(0)} > 0$ with high probability, we have $\langle \mathbf{w}_{j^{**}}^{(t)}, -\mathbf{v}\rangle > 0$. According to (A.6), it follows that

$$\begin{aligned}\underbrace{\sum_{i=1}^{n_{\mathrm{u}}} y_i\widehat{y}_i c_i^{(t)}[\langle \mathbf{w}_{j^{**}}^{(t)}, y_i\cdot\mathbf{v}\rangle]_+^2\|\mathbf{v}\|_2^2}_{\clubsuit} &= \sum_{i\in S_{-1}} y_i\widehat{y}_i c_i^{(t)}[\langle \mathbf{w}_{j^{**}}^{(t)}, -\mathbf{v}\rangle]_+^2\|\mathbf{v}\|_2^2 \\ &= n_{-1}\cdot\left(p-\frac{1}{2}\pm o(1)\right)\cdot\|\mathbf{v}\|_2^2\cdot\big(\widehat{\Lambda}_{-1}^{(t)}\big)^2,\end{aligned} \tag{A.12}$$

where $S_{-1} := \{(\mathbf{x}_i, y_i)|y_i = -1, i\in[n_{\mathrm{u}}]\}$, $n_{-1} = |S_{-1}|$.

For $\bigstar$, according to lemma A.7, similarly we have

$$\underbrace{\sum_{i=1}^{n_{\mathrm{l}}} b_i^{(t)}[\langle \mathbf{w}_{j^{**}}^{(t)}, y_i'\cdot\mathbf{v}\rangle]_+^2\|\mathbf{v}\|_2^2}_{\bigstar} = \sum_{i\in S'_{-1}} b_i^{(t)}[\langle \mathbf{w}_{j^{**}}^{(t)}, -\mathbf{v}\rangle]_+^2\|\mathbf{v}\|_2^2 = n'_{-1}\cdot\left(\frac{1}{2}\pm o(1)\right)\cdot\|\mathbf{v}\|_2^2\cdot\big(\widehat{\Lambda}_{-1}^{(t)}\big)^2, \tag{A.13}$$

where $S'_{-1} = \{(\mathbf{x}'_i, y'_i)|y'_i = -1, i\in[n_{\mathrm{l}}]\}$ and $n'_{-1} = |S'_{-1}|$.

According to (A.12) and (A.13), we have

$$\widehat{\Lambda}_{-1}^{(t+1)} \geq (1-\eta\lambda)\cdot\widehat{\Lambda}_{-1}^{(t)} + 3\eta\cdot\Bigg(\underbrace{\frac{n_{-1}}{n_{\mathrm{l}}+n_{\mathrm{u}}}\cdot\left(p-\frac{1}{2}\right) + \frac{n'_{-1}}{n_{\mathrm{l}}+n_{\mathrm{u}}}\cdot\frac{1}{2}}_{\spadesuit}\pm o(1)\Bigg)\cdot\|\mathbf{v}\|_2^2\cdot\big(\widehat{\Lambda}_{-1}^{(t)}\big)^2. \tag{A.14}$$

According to lemma A.9 and lemma A.10, and note that $n_{\mathrm{l}} = \widetilde{\Theta}(1), n_{\mathrm{u}} = \omega(d^{4\epsilon})$, we have for $\spadesuit$ that with probability at least $1 - 4\delta$

$$\begin{aligned}
&\Bigg| \underbrace{\frac{n_{-1}}{n_{\mathrm{l}} + n_{\mathrm{u}}} \cdot \Big(p - \frac{1}{2}\Big) + \frac{n'_{-1}}{n_{\mathrm{l}} + n_{\mathrm{u}}} \cdot \frac{1}{2}}_{\spadesuit} - \frac{n_{\mathrm{u}}}{2(n_{\mathrm{l}} + n_{\mathrm{u}})} \cdot \Big(p - \frac{1}{2}\Big) - \frac{n_{\mathrm{l}}}{2(n_{\mathrm{l}} + n_{\mathrm{u}})} \cdot \frac{1}{2} \Bigg| \\
&\leq \frac{|n_{-1} - \frac{n_{\mathrm{u}}}{2}|}{n_{\mathrm{l}} + n_{\mathrm{u}}} \cdot \Big(p - \frac{1}{2}\Big) + \frac{|n'_{-1} - \frac{n_{\mathrm{l}}}{2}|}{n_{\mathrm{l}} + n_{\mathrm{u}}} \cdot \frac{1}{2} \\
&\leq \frac{\sqrt{\frac{n_{\mathrm{u}}}{2} \log \frac{1}{\delta}}}{n_{\mathrm{l}} + n_{\mathrm{u}}} \cdot \Big(p - \frac{1}{2}\Big) + \frac{\sqrt{\frac{n_{\mathrm{l}}}{2} \log \frac{1}{\delta}}}{n_{\mathrm{l}} + n_{\mathrm{u}}} \cdot \frac{1}{2} \\
&= \Theta\Big(\frac{1}{\sqrt{n_{\mathrm{u}}}}\Big) \\
&= o(1).
\end{aligned}$$

Therefore, note that $n_{\mathrm{u}} = \omega(n_{\mathrm{l}})$ and $n_{\mathrm{u}} = \omega(1)$, we have

$$\begin{aligned}
\underbrace{\frac{n_{-1}}{n_{\mathrm{l}} + n_{\mathrm{u}}} \cdot \Big(p - \frac{1}{2}\Big) + \frac{n'_{-1}}{n_{\mathrm{l}} + n_{\mathrm{u}}} \cdot \frac{1}{2}}_{\spadesuit} &= \frac{n_{\mathrm{u}}}{2(n_{\mathrm{l}} + n_{\mathrm{u}})} \cdot \Big(p - \frac{1}{2}\Big) + \frac{n_{\mathrm{l}}}{2(n_{\mathrm{l}} + n_{\mathrm{u}})} \cdot \frac{1}{2} \pm o(1) \\
&= \frac{1}{2} \cdot \Big(p - \frac{1}{2}\Big) \pm o(1) \qquad \text{(A.15)}
\end{aligned}$$

Plugging (A.15) into (A.14), we have

$$\begin{aligned}
\widehat{\Lambda}_{-1}^{(t+1)} &\geq (1 - \eta\lambda) \cdot \widehat{\Lambda}_{-1}^{(t)} + 3\eta \cdot \bigg(\frac{1}{2} \cdot \Big(p - \frac{1}{2}\Big) \pm o(1)\bigg) \cdot \|\mathbf{v}\|_2^2 \cdot \big(\widehat{\Lambda}_{-1}^{(t)}\big)^2 \\
&= (1 - \eta\lambda) \cdot \widehat{\Lambda}_{-1}^{(t)} + \eta \cdot \Big(p - \frac{1}{2}\Big) \cdot \Theta(d) \cdot \big(\widehat{\Lambda}_{-1}^{(t)}\big)^2, \qquad \text{(A.16)}
\end{aligned}$$

which verifies the first inequality of case $r = -1$ in the lemma.

Next, we prove the latter part of the lemma. Let $j^\natural = \arg\max_{m+1 \leq j \leq 2m} \langle \mathbf{w}_j^{(t+1)}, \mathbf{v} \rangle$, then we have:

$$\begin{aligned}
\bar{\Lambda}_1^{(t+1)} &= \langle \mathbf{w}_{j^\natural}^{(t+1)}, \mathbf{v} \rangle \\
&= (1 - \eta\lambda) \cdot \langle \mathbf{w}_{j^\natural}^{(t)}, \mathbf{v} \rangle - \frac{3\eta}{n_{\mathrm{l}} + n_{\mathrm{u}}} \bigg(\underbrace{\sum_{i=1}^{n_{\mathrm{u}}} y_i \widehat{y}_i c_i^{(t)} [\langle \mathbf{w}_{j^\natural}^{(t)}, y_i \cdot \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2}_{\clubsuit} \\
&\qquad + \underbrace{\sum_{i=1}^{n_{\mathrm{l}}} b_i^{(t)} [\langle \mathbf{w}_{j^\natural}^{(t)}, y'_i \cdot \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2}_{\bigstar} \bigg).
\end{aligned}$$

For $\clubsuit$, according to (A.6), we have

$$\begin{aligned}
&\underbrace{\sum_{i=1}^{n_{\mathrm{u}}} y_i \widehat{y}_i c_i^{(t)} [\langle \mathbf{w}_{j^\natural}^{(t)}, y_i \cdot \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2}_{\clubsuit} \\
&= \sum_{i \in S_1} y_i \widehat{y}_i c_i^{(t)} [\langle \mathbf{w}_{j^\natural}^{(t)}, \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 + \sum_{i \in S_{-1}} y_i \widehat{y}_i c_i^{(t)} [\langle \mathbf{w}_{j^\natural}^{(t)}, -\mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 \\
&= \bigg(\sum_{i \in S_1} y_i \widehat{y}_i c_i^{(t)} \bigg) \cdot [\langle \mathbf{w}_{j^\natural}^{(t)}, \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 + \bigg(\sum_{i \in S_{-1}} y_i \widehat{y}_i c_i^{(t)} \bigg) \cdot [\langle \mathbf{w}_{j^\natural}^{(t)}, -\mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 \\
&= n_1 \cdot \Big(p - \frac{1}{2} \pm o(1)\Big) \cdot [\langle \mathbf{w}_{j^\natural}^{(t)}, \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 + n_{-1} \cdot \Big(p - \frac{1}{2} \pm o(1)\Big) \cdot [\langle \mathbf{w}_{j^\natural}^{(t)}, -\mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 \geq 0,
\end{aligned}$$

and for ★ it's obvious that

$$\underbrace{\sum_{i=1}^{n_l} b_i^{(t)} [\langle \mathbf{w}_{j^\natural}^{(t)}, y_i' \cdot \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2}_{\bigstar} \geq 0.$$

Therefore, it follows that

$$\bar{\Lambda}_1^{(t+1)} \leq (1-\eta\lambda) \cdot \langle \mathbf{w}_{j^\natural}^{(t)}, \mathbf{v} \rangle \leq (1-\eta\lambda) \bar{\Lambda}_1^{(t)}.$$

Let $j^{\natural\natural} = \arg\max_{1 \leq j \leq m} \langle \mathbf{w}_j^{(t+1)}, -\mathbf{v} \rangle$, then we have:

$$\begin{aligned}
\bar{\Lambda}_{-1}^{(t+1)} &= \langle \mathbf{w}_{j^{\natural\natural}}^{(t+1)}, -\mathbf{v} \rangle \\
&= (1-\eta\lambda) \cdot \langle \mathbf{w}_{j^{\natural\natural}}^{(t)}, -\mathbf{v} \rangle - \frac{3\eta}{n_l + n_u} \Bigg(\sum_{i=1}^{n_u} y_i \widehat{y}_i c_i^{(t)} [\langle \mathbf{w}_{j^{\natural\natural}}^{(t)}, y_i \cdot \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 \\
&\qquad + \sum_{i=1}^{n_l} b_i^{(t)} [\langle \mathbf{w}_{j^{\natural\natural}}^{(t)}, y_i' \cdot \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 \Bigg) \\
&\leq (1-\eta\lambda) \cdot \langle \mathbf{w}_{j^{\natural\natural}}^{(t)}, -\mathbf{v} \rangle \\
&\leq (1-\eta\lambda) \cdot \bar{\Lambda}_{-1}^{(t)},
\end{aligned}$$

which verifies the second part of the lemma. □

Although the accuracy of pseudolabeler is larger than $1/2$, which is used as an assumption in the previous proof, we can also analyse the model with high label flipping probability and the accuracy of pseudolabeler p is smaller than $1/2$. In this case, the neural network will turn to fit the opposite direction of feature vector, $\bar{\Lambda}_r^{(t)}$ will increase and $\widehat{\Lambda}_r^{(t)}$ will decrease, which is formulated as the following lemma.

Lemma A.12. For $\widehat{\Lambda}_1^{(t)} := \max_{1 \leq j \leq m} \langle \mathbf{w}_j^{(t)}, \mathbf{v} \rangle$ and $\widehat{\Lambda}_{-1}^{(t)} := \max_{m+1 \leq j \leq 2m} \langle \mathbf{w}_j^{(t)}, -\mathbf{v} \rangle$, we have with high probability that

$$\widehat{\Lambda}_r^{(t+1)} \leq (1-\eta\lambda) \cdot \widehat{\Lambda}_r^{(t)}, r \in \{\pm 1\}.$$

For $\bar{\Lambda}_1^{(t)} := \max_{m+1 \leq j \leq 2m} \langle \mathbf{w}_j^{(t)}, \mathbf{v} \rangle$ and $\bar{\Lambda}_1^{(t)} := \max_{1 \leq j \leq m} \langle \mathbf{w}_j^{(t)}, -\mathbf{v} \rangle$, we have with high probability that

$$\bar{\Lambda}_r^{(t+1)} \geq (1-\eta\lambda) \cdot \bar{\Lambda}_r^{(t)} + \eta \cdot \left(\frac{1}{2} - p \right) \cdot \Theta(d) \cdot (\bar{\Lambda}_r^{(t)})^2, r \in \{\pm 1\}.$$

Proof of lemma A.12. First, we prove the former part of this lemma. Let $j^* = \arg\max_{1 \leq j \leq m} \langle \mathbf{w}_j^{(t+1)}, \mathbf{v} \rangle$ and note that $u_{j^*} = \mathbb{1}_{[1 \leq j \leq m]} - \mathbb{1}_{[m+1 \leq j \leq 2m]} = 1$, then we have

$$\begin{aligned}
\widehat{\Lambda}_1^{(t+1)} &= \langle \mathbf{w}_{j^*}^{(t+1)}, \mathbf{v} \rangle \\
&= (1-\eta\lambda) \cdot \langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v} \rangle + \frac{3\eta}{n_l + n_u} \Bigg(\underbrace{\sum_{i=1}^{n_u} y_i \widehat{y}_i c_i^{(t)} [\langle \mathbf{w}_{j^*}^{(t)}, y_i \cdot \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2}_{\clubsuit} \\
&\qquad + \underbrace{\sum_{i=1}^{n_l} b_i^{(t)} [\langle \mathbf{w}_{j^*}^{(t)}, y_i' \cdot \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2}_{\bigstar} \Bigg).
\end{aligned}$$

For ♣, according to (A.6), we have

$$
\begin{aligned}
&\underbrace{\sum_{i=1}^{n_{\mathrm{u}}} y_i \widehat{y}_i c_i^{(t)} [\langle \mathbf{w}_{j^*}^{(t)}, y_i \cdot \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2}_{\clubsuit} \\
&= \sum_{i \in S_1} y_i \widehat{y}_i c_i^{(t)} [\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 + \sum_{i \in S_{-1}} y_i \widehat{y}_i c_i^{(t)} [\langle \mathbf{w}_{j^*}^{(t)}, -\mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 \\
&= \bigg(\sum_{i \in S_1} y_i \widehat{y}_i c_i^{(t)} \bigg) \cdot [\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 + \bigg(\sum_{i \in S_{-1}} y_i \widehat{y}_i c_i^{(t)} \bigg) \cdot [\langle \mathbf{w}_{j^*}^{(t)}, -\mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 \\
&= n_1 \cdot \Big(p - \frac{1}{2} \pm o(1) \Big) \cdot [\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 + n_{-1} \cdot \Big(p - \frac{1}{2} \pm o(1) \Big) \cdot [\langle \mathbf{w}_{j^*}^{(t)}, -\mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2,
\end{aligned}
$$

For ★, according to (A.6), we have

$$
\begin{aligned}
&\underbrace{\sum_{i=1}^{n_{\mathrm{l}}} b_i^{(t)} [\langle \mathbf{w}_{j^*}^{(t)}, y_i' \cdot \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2}_{\bigstar} \\
&= \sum_{i \in S_1'} b_i^{(t)} [\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 + \sum_{i \in S_{-1}'} b_i^{(t)} [\langle \mathbf{w}_{j^*}^{(t)}, -\mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 \\
&= n_1' \cdot \Big(\frac{1}{2} \pm o(1) \Big) \cdot [\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 + n_{-1}' \cdot \Big(\frac{1}{2} \pm o(1) \Big) \cdot [\langle \mathbf{w}_{j^*}^{(t)}, -\mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2,
\end{aligned}
$$

It follows that

$$
\begin{aligned}
&\underbrace{\sum_{i=1}^{n_{\mathrm{u}}} y_i \widehat{y}_i c_i^{(t)} [\langle \mathbf{w}_{j^*}^{(t)}, y_i \cdot \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2}_{\clubsuit} + \underbrace{\sum_{i=1}^{n_{\mathrm{l}}} b_i^{(t)} [\langle \mathbf{w}_{j^*}^{(t)}, y_i' \cdot \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2}_{\bigstar} \\
&= \bigg(n_1 \cdot \Big(p - \frac{1}{2} \pm o(1) \Big) + n_1' \cdot \Big(\frac{1}{2} \pm o(1) \Big) \bigg) \cdot [\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 \\
&\quad + \bigg(n_{-1} \cdot \Big(p - \frac{1}{2} \pm o(1) \Big) + n_{-1}' \cdot \Big(\frac{1}{2} \pm o(1) \Big) \bigg) \cdot [\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2.
\end{aligned}
$$

According to lemma A.9 and note that $n_{\mathrm{u}} = \omega(n_{\mathrm{l}})$, it holds with probability at least $1 - 8\delta$ that

$$
\begin{aligned}
n_1' \cdot \Big(\frac{1}{2} \pm o(1) \Big) &\leq \bigg(\frac{n_{\mathrm{l}}}{2} + \sqrt{\frac{n_{\mathrm{l}}}{2} \log \frac{1}{\delta}} \bigg) \cdot \Big(\frac{1}{2} \pm o(1) \Big) = \Theta(n_{\mathrm{l}}) = o(n_{\mathrm{u}}) \\
&\leq \bigg(\frac{n_{\mathrm{u}}}{2} + \sqrt{\frac{n_{\mathrm{u}}}{2} \log \frac{1}{\delta}} \bigg) \cdot \Big(\frac{1}{2} - p \pm o(1) \Big) \leq n_1 \cdot \Big(\frac{1}{2} - p \pm o(1) \Big),
\end{aligned}
$$

$$
\begin{aligned}
n_{-1}' \cdot \Big(\frac{1}{2} \pm o(1) \Big) &\leq \bigg(\frac{n_{\mathrm{l}}}{2} + \sqrt{\frac{n_{\mathrm{l}}}{2} \log \frac{1}{\delta}} \bigg) \cdot \Big(\frac{1}{2} \pm o(1) \Big) = \Theta(n_{\mathrm{l}}) = o(n_{\mathrm{u}}) \\
&\leq \bigg(\frac{n_{\mathrm{u}}}{2} + \sqrt{\frac{n_{\mathrm{u}}}{2} \log \frac{1}{\delta}} \bigg) \cdot \Big(\frac{1}{2} - p \pm o(1) \Big) \leq n_{-1} \cdot \Big(\frac{1}{2} - p \pm o(1) \Big),
\end{aligned}
$$

leading to $\clubsuit + \bigstar \leq 0$. Therefore,

$$
\widehat{\Lambda}_1^{(t+1)} \leq (1 - \eta\lambda) \langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v} \rangle \leq (1 - \eta\lambda) \cdot \widehat{\Lambda}_1^{(t)}.
$$

And we can prove in a similar way that $\widehat{\Lambda}_{-1}^{(t+1)} \leq (1 - \eta\lambda) \cdot \widehat{\Lambda}_{-1}^{(t)}$.

Next, we prove the second part of the lemma. Let $j^\natural = \arg\max_{m+1\le j\le 2m}\langle \mathbf{w}_j^{(t)}, \mathbf{v}\rangle$ and note that $u_{j^\natural} = \mathbb{1}_{[1\le j\le m]} - \mathbb{1}_{[m+1\le j\le 2m]} = -1$, then we have

$$\begin{aligned}
\bar{\Lambda}_1^{(t+1)} &\ge \langle \mathbf{w}_{j^\natural}^{(t+1)}, \mathbf{v}\rangle \\
&= (1-\eta\lambda)\cdot\langle \mathbf{w}_{j^\natural}^{(t)}, \mathbf{v}\rangle - \frac{3\eta}{n_{\mathrm{l}}+n_{\mathrm{u}}}\Bigg(\underbrace{\sum_{i=1}^{n_{\mathrm{u}}} y_i\widehat{y}_i c_i^{(t)}[\langle \mathbf{w}_{j^\natural}^{(t)}, y_i\cdot\mathbf{v}\rangle]_+^2\|\mathbf{v}\|_2^2}_{\clubsuit} \\
&\quad + \underbrace{\sum_{i=1}^{n_{\mathrm{l}}} b_i^{(t)}[\langle \mathbf{w}_{j^\natural}^{(t)}, y_i'\cdot\mathbf{v}\rangle]_+^2\|\mathbf{v}\|_2^2}_{\bigstar}\Bigg).
\end{aligned}$$

For $\clubsuit$, note the definition of $j^\natural$ that $\bar{\Lambda}_1^{(t)} = \langle \mathbf{w}_{j^\natural}^{(t)}, \mathbf{v}\rangle$ and note the increasing property of $\bar{\Lambda}_1^{(t)}$ in this case and $\bar{\Lambda}_1^{(0)} > 0$ with high probability, we have $\langle \mathbf{w}_{j^\natural}^{(t)}, \mathbf{v}\rangle > 0$. It follows that

$$\begin{aligned}
\underbrace{\sum_{i=1}^{n_{\mathrm{u}}} y_i\widehat{y}_i c_i^{(t)}[\langle \mathbf{w}_{j^\natural}^{(t)}, y_i\cdot\mathbf{v}\rangle]_+^2\|\mathbf{v}\|_2^2}_{\clubsuit} &= \sum_{i\in S_1} y_i\widehat{y}_i c_i^{(t)}[\langle \mathbf{w}_{j^\natural}^{(t)}, \mathbf{v}\rangle]_+^2\|\mathbf{v}\|_2^2 + \sum_{i\in S_{-1}} y_i\widehat{y}_i c_i^{(t)}[-\langle \mathbf{w}_{j^\natural}^{(t)}, \mathbf{v}\rangle]_+^2\|\mathbf{v}\|_2^2 \\
&= \sum_{i\in S_1} y_i\widehat{y}_i c_i^{(t)}[\langle \mathbf{w}_{j^\natural}^{(t)}, \mathbf{v}\rangle]_+^2\|\mathbf{v}\|_2^2 \\
&= \Bigg(\sum_{i\in S_1} y_i\widehat{y}_i c_i^{(t)}\Bigg)\cdot\|\mathbf{v}\|_2^2\cdot\big(\bar{\Lambda}_1^{(t)}\big)^2 \\
&= n_1\cdot\Big(p-\frac{1}{2}\pm o(1)\Big)\cdot\|\mathbf{v}\|_2^2\cdot\big(\bar{\Lambda}_1^{(t)}\big)^2, \qquad \text{(A.17)}
\end{aligned}$$

For $\bigstar$, similarly we have

$$\begin{aligned}
\underbrace{\sum_{i=1}^{n_{\mathrm{l}}} b_i^{(t)}[\langle \mathbf{w}_{j^\natural}^{(t)}, y_i'\cdot\mathbf{v}\rangle]_+^2\|\mathbf{v}\|_2^2}_{\bigstar} &= \sum_{i\in S_1'} b_i^{(t)}[\langle \mathbf{w}_{j^\natural}^{(t)}, \mathbf{v}\rangle]_+^2\|\mathbf{v}\|_2^2 + \sum_{i\in S_{-1}'} b_i^{(t)}[-\langle \mathbf{w}_{j^\natural}^{(t)}, \mathbf{v}\rangle]_+^2\|\mathbf{v}\|_2^2 \\
&= \sum_{i\in S_1'} b_i^{(t)}[\langle \mathbf{w}_{j^\natural}^{(t)}, \mathbf{v}\rangle]_+^2\|\mathbf{v}\|_2^2 \\
&= \Bigg(\sum_{i\in S_1'} b_i^{(t)}\Bigg)\cdot\|\mathbf{v}\|_2^2\cdot\big(\bar{\Lambda}_1^{(t)}\big)^2 \\
&= n_1'\cdot\Big(\frac{1}{2}\pm o(1)\Big)\cdot\|\mathbf{v}\|_2^2\cdot\big(\bar{\Lambda}_1^{(t)}\big)^2. \qquad \text{(A.18)}
\end{aligned}$$

According to lemma A.9, (A.17) and (A.18), we have $n_1' = o(n_1)$ with high probability, therefore

$$\clubsuit + \bigstar = n_1\cdot\Big(p-\frac{1}{2}\pm o(1)\Big)\cdot\|\mathbf{v}\|_2^2\cdot\big(\bar{\Lambda}_1^{(t)}\big)^2,$$

leading to

$$\begin{aligned}
\bar{\Lambda}_1^{(t+1)} &\ge (1-\eta\lambda)\cdot\langle \mathbf{w}_{j^\natural}^{(t)}, \mathbf{v}\rangle - \frac{3\eta n_1}{n_{\mathrm{l}}+n_{\mathrm{u}}}\cdot\Big(p-\frac{1}{2}\pm o(1)\Big)\cdot\|\mathbf{v}\|_2^2\cdot\big(\bar{\Lambda}_1^{(t)}\big)^2 \\
&= (1-\eta\lambda)\cdot\bar{\Lambda}_1^{(t)} + \eta\cdot\Big(\frac{1}{2}-p\Big)\cdot\Theta(d)\cdot\big(\bar{\Lambda}_1^{(t)}\big)^2.
\end{aligned}$$

And we can prove in a similar way that

$$\bar{\Lambda}_1^{(t+1)} \ge (1-\eta\lambda)\cdot\bar{\Lambda}_1^{(t)} + \eta\cdot\Big(\frac{1}{2}-p\Big)\cdot\Theta(d)\cdot\big(\bar{\Lambda}_1^{(t)}\big)^2.$$

□

A.4.2 UNIFORM UPPER BOUND FOR $\Gamma^{(t)}$

The following lemma provides an upper bound for the increasing rate of $\Gamma^{(t)}$.

Lemma A.13. For $\Gamma_i^{(t)} := \max_{j\in[2m]}\langle \mathbf{w}_j, \boldsymbol{\xi}_i\rangle, i \in [n_{\mathrm{u}}]$, $\Gamma_i'^{(t)} := \max_{j\in[2m]}\langle \mathbf{w}_j, \boldsymbol{\xi}_i'\rangle, i \in [n_{\mathrm{l}}]$, $\Gamma^{(t)} := \max\{\max_{i\in[n_{\mathrm{u}}]}\Gamma_i^{(t)}, \max_{i\in[n_{\mathrm{l}}]}\Gamma_i'^{(t)}\}$, we have with high probability that

$$\Gamma_i^{(t+1)} \le (1-\eta\lambda)\cdot\Gamma_i^{(t)} + \eta\cdot\max\left\{\widetilde{\Theta}(d^{\frac{1}{2}+2\epsilon}), \widetilde{\Theta}\Big(\frac{d^{1+2\epsilon}}{n_{\mathrm{u}}}\Big)\right\}\cdot\big(\Gamma^{(t)}\big)^2, i\in[n_{\mathrm{l}}],$$

$$\Gamma_i'^{(t+1)} \le (1-\eta\lambda)\cdot\Gamma_i'^{(t)} + \eta\cdot\max\left\{\widetilde{\Theta}(d^{\frac{1}{2}+2\epsilon}), \widetilde{\Theta}\Big(\frac{d^{1+2\epsilon}}{n_{\mathrm{u}}}\Big)\right\}\cdot\big(\Gamma^{(t)}\big)^2, i\in[n_{\mathrm{l}}],$$

and

$$\Gamma^{(t+1)} \le (1-\eta\lambda)\cdot\Gamma^{(t)} + \eta\cdot\max\left\{\widetilde{\Theta}(d^{\frac{1}{2}+2\epsilon}), \widetilde{\Theta}\Big(\frac{d^{1+2\epsilon}}{n_{\mathrm{u}}}\Big)\right\}\cdot\big(\Gamma^{(t)}\big)^2.$$

Proof of lemma A.13. We first prove the former inequality. Let $j^\star = \arg\max_{1\le j\le 2m}\langle \mathbf{w}_j^{(t+1)}, \boldsymbol{\xi}_l\rangle$, where $l\in[n_{\mathrm{u}}]$ is fixed. According to lemma A.2, we have

$$\begin{aligned}
&\Gamma_l^{(t+1)} = \langle \mathbf{w}_{j^\star}^{(t+1)}, \boldsymbol{\xi}_l\rangle \\
&= (1-\eta\lambda)\cdot\langle \mathbf{w}_{j^\star}^{(t)}, \boldsymbol{\xi}_l\rangle + \frac{3\eta u_{j^\star}}{n_{\mathrm{l}}+n_{\mathrm{u}}}\bigg(\sum_{i=1}^{n_{\mathrm{u}}}\widehat{y}_i c_i^{(t)}[\langle \mathbf{w}_{j^\star}^{(t)}, \boldsymbol{\xi}_i\rangle]_+^2\langle\boldsymbol{\xi}_i,\boldsymbol{\xi}_l\rangle + \sum_{i=1}^{n_{\mathrm{l}}} y_i' b_i^{(t)}[\langle \mathbf{w}_{j^\star}^{(t)}, \boldsymbol{\xi}_i'\rangle]_+^2\langle\boldsymbol{\xi}_i',\boldsymbol{\xi}_l\rangle\bigg) \\
&\le (1-\eta\lambda)\cdot\langle \mathbf{w}_{j^\star}^{(t)}, \boldsymbol{\xi}_l\rangle + \frac{3\eta}{n_{\mathrm{l}}+n_{\mathrm{u}}}\bigg(\underbrace{\sum_{i=1}^{n_{\mathrm{u}}} c_i^{(t)}[\langle \mathbf{w}_{j^\star}^{(t)}, \boldsymbol{\xi}_i\rangle]_+^2|\langle\boldsymbol{\xi}_i,\boldsymbol{\xi}_l\rangle|}_{\clubsuit} + \underbrace{\sum_{i=1}^{n_{\mathrm{l}}} b_i^{(t)}[\langle \mathbf{w}_{j^\star}^{(t)}, \boldsymbol{\xi}_i'\rangle]_+^2|\langle\boldsymbol{\xi}_i',\boldsymbol{\xi}_l\rangle|}_{\bigstar}\bigg),
\end{aligned} \tag{A.19}$$

where the last inequality is due to triangle inequality.

For $\clubsuit$, note that $l\in[n_{\mathrm{u}}]$ and there exists an $i\in[n_{\mathrm{u}}]$ equivalent to l, it follows that

$$\begin{aligned}
&\underbrace{\sum_{i=1}^{n_{\mathrm{u}}} c_i^{(t)}[\langle \mathbf{w}_{j^\star}^{(t)}, \boldsymbol{\xi}_i\rangle]_+^2|\langle\boldsymbol{\xi}_i,\boldsymbol{\xi}_l\rangle|}_{\clubsuit} \\
&= \sum_{i\in[n_{\mathrm{u}}], i\ne l} c_i^{(t)}[\langle \mathbf{w}_{j^\star}^{(t)}, \boldsymbol{\xi}_i\rangle]_+^2|\langle\boldsymbol{\xi}_i,\boldsymbol{\xi}_l\rangle| + c_l^{(t)}[\langle \mathbf{w}_{j^\star}^{(t)}, \boldsymbol{\xi}_l\rangle]_+^2\|\boldsymbol{\xi}_l\|_2^2 \\
&\le (n_{\mathrm{u}}-1)\cdot\Big(\frac{1}{2}+o(1)\Big)\cdot\widetilde{\Theta}(d^{\frac{1}{2}+2\epsilon})\cdot\big(\Gamma^{(t)}\big)^2 + \Big(\frac{1}{2}+o(1)\Big)\cdot\widetilde{\Theta}(d^{1+2\epsilon})\cdot\big(\Gamma^{(t)}\big)^2 \\
&= (n_{\mathrm{u}}-1)\cdot\widetilde{\Theta}(d^{\frac{1}{2}+2\epsilon})\cdot\big(\Gamma^{(t)}\big)^2 + \widetilde{\Theta}(d^{1+2\epsilon})\cdot\big(\Gamma^{(t)}\big)^2,
\end{aligned} \tag{A.20}$$

where the inequality is due to lemma A.7, $\|\boldsymbol{\xi}_l\|_2^2 = \widetilde{\Theta}(d\sigma_p^2) = \widetilde{\Theta}(d^{1+2\epsilon})$, $|\langle\boldsymbol{\xi}_i,\boldsymbol{\xi}_l\rangle| = \widetilde{\Theta}(d^{\frac{1}{2}}\sigma_p^2) = \widetilde{\Theta}(d^{\frac{1}{2}+2\epsilon})$ according to Lemma C.1 and the definition of $\Gamma^{(t)}$.

For $\bigstar$, we have

$$\underbrace{\sum_{i=1}^{n_{\mathrm{l}}} b_i^{(t)}[\langle \mathbf{w}_{j^\star}^{(t)}, \boldsymbol{\xi}_i'\rangle]_+^2|\langle\boldsymbol{\xi}_i',\boldsymbol{\xi}_l\rangle|}_{\bigstar} \le n_{\mathrm{l}}\cdot\Big(\frac{1}{2}+o(1)\Big)\cdot\widetilde{\Theta}(d^{\frac{1}{2}+2\epsilon})\cdot\big(\Gamma^{(t)}\big)^2 = n_{\mathrm{l}}\cdot\widetilde{\Theta}(d^{\frac{1}{2}+2\epsilon})\cdot\big(\Gamma^{(t)}\big)^2, \tag{A.21}$$

Plugging (A.20) and (A.21) into (A.19), we have

$$\begin{aligned}
\Gamma_l^{(t+1)} &\le (1-\eta\lambda)\cdot\Gamma_l^{(t+1)} + \eta\cdot\bigg(\frac{3}{n_{\mathrm{l}}+n_{\mathrm{u}}}\cdot\Big((n_{\mathrm{u}}+n_{\mathrm{l}}-1)\cdot\widetilde{\Theta}(d^{\frac{1}{2}+2\epsilon}) + \widetilde{\Theta}(d^{1+2\epsilon})\Big)\bigg)\cdot\big(\Gamma^{(t)}\big)^2 \\
&\le (1-\eta\lambda)\cdot\Gamma_l^{(t+1)} + \eta\cdot\max\left\{\widetilde{\Theta}(d^{\frac{1}{2}+2\epsilon}), \widetilde{\Theta}\Big(\frac{d^{1+2\epsilon}}{n_{\mathrm{u}}}\Big)\right\}\cdot\big(\Gamma^{(t)}\big)^2,
\end{aligned}$$

which is the first part of this lemma.

Let $j^\star = \mathrm{argmax}_{1\le j\le 2m}\langle \mathbf{w}_j^{(t+1)}, \boldsymbol{\xi}_l'\rangle$, where $l\in[n_{\mathrm{l}}]$ is fixed. According to lemma A.2, we have

$$\begin{aligned}
\Gamma_l'^{(t+1)} &= \langle \mathbf{w}_{j^\star}^{(t+1)}, \boldsymbol{\xi}_l'\rangle \\
&= (1-\eta\lambda)\cdot\langle \mathbf{w}_{j^\star}^{(t)}, \boldsymbol{\xi}_l'\rangle + \frac{3\eta u_{j^\star}}{n_{\mathrm{l}}+n_{\mathrm{u}}}\bigg(\sum_{i=1}^{n_{\mathrm{u}}}\widehat{y}_i c_i^{(t)}[\langle \mathbf{w}_{j^\star}^{(t)}, \boldsymbol{\xi}_i\rangle]_+^2\langle \boldsymbol{\xi}_i, \boldsymbol{\xi}_l'\rangle + \sum_{i=1}^{n_{\mathrm{l}}} y_i' b_i^{(t)}[\langle \mathbf{w}_{j^\star}^{(t)}, \boldsymbol{\xi}_i'\rangle]_+^2\langle \boldsymbol{\xi}_i', \boldsymbol{\xi}_l'\rangle\bigg) \\
&\le (1-\eta\lambda)\cdot\langle \mathbf{w}_{j^\star}^{(t)}, \boldsymbol{\xi}_l'\rangle + \frac{3\eta}{n_{\mathrm{l}}+n_{\mathrm{u}}}\bigg(\underbrace{\sum_{i=1}^{n_{\mathrm{u}}} c_i^{(t)}[\langle \mathbf{w}_{j^\star}^{(t)}, \boldsymbol{\xi}_i\rangle]_+^2|\langle \boldsymbol{\xi}_i, \boldsymbol{\xi}_l'\rangle|}_{\clubsuit} + \underbrace{\sum_{i=1}^{n_{\mathrm{l}}} b_i^{(t)}[\langle \mathbf{w}_{j^\star}^{(t)}, \boldsymbol{\xi}_i'\rangle]_+^2|\langle \boldsymbol{\xi}_i', \boldsymbol{\xi}_l'\rangle|}_{\bigstar}\bigg),
\end{aligned} \tag{A.22}$$

For $\clubsuit$, we have

$$\underbrace{\sum_{i=1}^{n_{\mathrm{u}}} c_i^{(t)}[\langle \mathbf{w}_{j^\star}^{(t)}, \boldsymbol{\xi}_i\rangle]_+^2|\langle \boldsymbol{\xi}_i, \boldsymbol{\xi}_l\rangle|}_{\clubsuit} \le \sum_{i=1}^{n_{\mathrm{u}}}\Big(\frac{1}{2}\pm o(1)\Big)\cdot\widetilde{\Theta}(d^{\frac{1}{2}+2\epsilon})\cdot\big(\Gamma^{(t)}\big)^2 = n_{\mathrm{u}}\cdot\widetilde{\Theta}(d^{\frac{1}{2}+2\epsilon})\cdot\big(\Gamma^{(t)}\big)^2, \tag{A.23}$$

where the inequality is due to lemma A.7, $|\langle \boldsymbol{\xi}_i, \boldsymbol{\xi}_l\rangle| = \widetilde{\Theta}(d^{\frac{1}{2}}\sigma_p^2) = \widetilde{\Theta}(d^{\frac{1}{2}+2\epsilon})$ and the definition of $\Gamma^{(t)}$.

For $\bigstar$, note that $l\in[n_{\mathrm{l}}]$ and there exists an $i\in[n_{\mathrm{l}}]$ equivalent to l, it follows that

$$\begin{aligned}
&\underbrace{\sum_{i=1}^{n_{\mathrm{l}}} b_i^{(t)}[\langle \mathbf{w}_{j^\star}^{(t)}, \boldsymbol{\xi}_i'\rangle]_+^2|\langle \boldsymbol{\xi}_i', \boldsymbol{\xi}_l'\rangle|}_{\bigstar} \\
&= \sum_{i\in[n_{\mathrm{l}}], i\ne l} b_i^{(t)}[\langle \mathbf{w}_{j^\star}^{(t)}, \boldsymbol{\xi}_i'\rangle]_+^2|\langle \boldsymbol{\xi}_i', \boldsymbol{\xi}_l'\rangle| + b_l^{(t)}[\langle \mathbf{w}_{j^\star}^{(t)}, \boldsymbol{\xi}_l'\rangle]_+^2\|\boldsymbol{\xi}_l'\|_2^2 \\
&\le (n_{\mathrm{l}}-1)\cdot\Big(\frac{1}{2}+o(1)\Big)\cdot\widetilde{\Theta}(d^{\frac{1}{2}+2\epsilon})\cdot\big(\Gamma^{(t)}\big)^2 + \Big(\frac{1}{2}+o(1)\Big)\cdot\widetilde{\Theta}(d^{1+2\epsilon})\cdot\big(\Gamma^{(t)}\big)^2 \\
&= (n_{\mathrm{l}}-1)\cdot\widetilde{\Theta}(d^{\frac{1}{2}+2\epsilon}) + \widetilde{\Theta}(d^{1+2\epsilon})\cdot\big(\Gamma^{(t)}\big)^2
\end{aligned} \tag{A.24}$$

Plugging (A.23) and (A.24) into (A.22), we have

$$\begin{aligned}
\Gamma_l'^{(t+1)} &\le (1-\eta\lambda)\cdot\Gamma_l'^{(t+1)} + \eta\cdot\bigg(\frac{3}{n_{\mathrm{l}}+n_{\mathrm{u}}}\cdot\Big((n_{\mathrm{u}}+n_{\mathrm{l}}-1)\cdot\widetilde{\Theta}(d^{\frac{1}{2}+2\epsilon}) + \widetilde{\Theta}(d^{1+2\epsilon})\Big)\bigg)\cdot\big(\Gamma^{(t)}\big)^2 \\
&\le (1-\eta\lambda)\cdot\Gamma_l'^{(t+1)} + \eta\cdot\max\Big\{\widetilde{\Theta}(d^{\frac{1}{2}+2\epsilon}), \widetilde{\Theta}\Big(\frac{d^{1+2\epsilon}}{n_{\mathrm{u}}}\Big)\Big\}\cdot\big(\Gamma^{(t)}\big)^2,
\end{aligned}$$

which verifies the second inequality in this lemma.

Note that $\Gamma^{(t)} = \max\{\max_{l\in[n_{\mathrm{u}}]}\Gamma_l^{(t)}, \max_{l\in[n_{\mathrm{l}}]}\Gamma_l'^{(t)}\}$, without loss of generality, we assume $\Gamma^{(t)} = \max_{l\in[n_{\mathrm{u}}]}\Gamma_l^{(t)}$ and assume $l^* = \mathrm{argmax}_{l\in[n_{\mathrm{u}}]}\Gamma_l^{(t+1)}$, we have

$$\begin{aligned}
\Gamma^{(t+1)} = \Gamma_{l^*}^{(t+1)} &\le (1-\eta\lambda)\cdot\Gamma_{l^*}^{(t)} + \eta\cdot\max\Big\{\widetilde{\Theta}(d^{\frac{1}{2}+2\epsilon}), \widetilde{\Theta}\Big(\frac{d^{1+2\epsilon}}{n_{\mathrm{u}}}\Big)\Big\}\cdot\big(\Gamma^{(t)}\big)^2 \\
&\le (1-\eta\lambda)\cdot\Gamma^{(t)} + \eta\cdot\max\Big\{\widetilde{\Theta}(d^{\frac{1}{2}+2\epsilon}), \widetilde{\Theta}\Big(\frac{d^{1+2\epsilon}}{n_{\mathrm{u}}}\Big)\Big\}\cdot\big(\Gamma^{(t)}\big)^2,
\end{aligned}$$

which verifies the third inequality in this lemma.

□

A.4.3 PROVE OFF-DIAGONAL CORRELATION REMAINS INITIALIZATION MAGNITUDE DURING $[0, T_r]$

If the accuracy of pseudolabeler $p > 1/2$, we have off-diagonal correlation $\bar{\Lambda}_r^{(t+1)} \leq (1-\eta\lambda)\cdot\bar{\Lambda}_r^{(t)}$ for $r \in \{+1,-1\}$, therefore, $\bar{\Lambda}_r^{(t)} = O(\bar{\Lambda}_r^{(0)}) = \widetilde{O}(d^{-\frac{1}{4}})$. If $p < 1/2$, we have off-diagonal correlation $\widehat{\Lambda}_r^{(t+1)} \leq (1-\eta\lambda)\cdot\widehat{\Lambda}_r^{(t)}$ for $r \in \{+1,-1\}$, therefore, $\widehat{\Lambda}_r^{(t)} = O(\widehat{\Lambda}_r^{(0)}) = \widetilde{O}(d^{-\frac{1}{4}})$.

A.4.4 TENSOR POWER METHOD: PROVING $\Gamma^{(t)} = O(\Gamma^{(0)})$ DURING $[0, T_r]$ AND COMPUTING THE MAGNITUDE OF T_r

According to section A.4.1 and A.4.2, we have obtained following upper bounds and lower bounds for *feature learning* term $\widehat{\Lambda}_r^{(t)}, \bar{\Lambda}_r^{(t)}, r \in \{+1,-1\}$ and *noise memorization* term $\Gamma^{(t)}$: When $t \in [0, T_r]$, we have

$$\begin{aligned}
&\widehat{\Lambda}_r^{(t+1)} \geq \widehat{\Lambda}_r^{(t)} + \eta\cdot(2p-1)\cdot\Theta(d)\cdot(\widehat{\Lambda}_r^{(t)})^2 \text{ and } \bar{\Lambda}_r^{(t+1)} \leq (1-\eta\lambda)\cdot\bar{\Lambda}_r^{(t)}, \text{ for } r\in\{+1,-1\};\\
&\Gamma^{(t+1)} \leq (1-\eta\lambda)\cdot\Gamma^{(t)} + \eta\cdot\max\left\{\widetilde{\Theta}(d^{\frac{1}{2}+2\epsilon}), \widetilde{\Theta}\Big(\frac{d^{1+2\epsilon}}{n_{\mathrm{u}}}\Big)\right\}\cdot(\Gamma^{(t)})^2.
\end{aligned} \tag{A.25}$$

According to condition 4.1, assume $n_{\mathrm{u}} = \Omega(d^{4\epsilon})$, we have

$$\max\left\{\widetilde{\Theta}(d^{\frac{1}{2}+2\epsilon}), \widetilde{\Theta}\Big(\frac{d^{1+2\epsilon}}{n_{\mathrm{u}}}\Big)\right\} = \widetilde{\Omega}(d^{1-2\epsilon}),$$

leading to

$$\Gamma^{(t+1)} \leq (1-\eta\lambda)\cdot\Gamma^{(t)} + \eta\cdot\widetilde{\Theta}(d^{1-2\epsilon})\cdot(\Gamma^{(t)})^2.$$

By leveraging tensor power method introduced in Lemma C.2, we can prove following lemma about the magnitude of $\Gamma^{(t)}$:

Lemma A.14. $\Gamma^{(t)}$ remains initialization magnitude during $[0, \max_{r\in\{\pm1\}}\{T_r\}]$.

Proof of lemma A.14. Let T_r^* be the first iteration t in which $\widehat{\Lambda}_r^{(t)} \geq A$ for $r \in \{+1,-1\}$, let T^* be the first iteration t in which $\Gamma^{(t)} \geq A'$, then according to lemma C.2, we know

$$\eta\cdot T_r^*\cdot C_t = \sum_{t\geq 0, \Lambda_r^{(t)}\leq A} \eta\cdot C_t \leq \frac{1+\delta}{\widehat{\Lambda}_r^{(0)}} + \eta(1+\delta)^2\Big(1+\frac{\log(A/\widehat{\Lambda}_r^{(0)})}{\log(1+\delta)}\Big);$$

$$\eta\cdot T^*\cdot C_t' = \sum_{t\geq 0, \Gamma^{(t)}\leq A'} \eta\cdot C_t' \geq \frac{1-\frac{x_0}{A'}}{(1+\delta')\Gamma^{(0)}} - \eta(1+\delta')^{-2}\Big(1+\frac{\log(A'/\Gamma^{(0)})}{\log(1+\delta')}\Big)$$

where $C_t = (2p-1)\cdot\widetilde{\Theta}(d)$ and $C_t' = \widetilde{\Theta}(d^{1-2\epsilon})$ according to (A.25).

Taking $A = \Theta(1/m), A' = C\cdot\Gamma^{(t)}$ where C is a large constant and $C = \Theta(1)$, $\delta = \delta' = \frac{1}{2}$ and note that $\widehat{\Lambda}_r^{(0)} = \widetilde{\Theta}(\sigma_0 d^{\frac{1}{2}}) = \widetilde{\Theta}(d^{-\frac{1}{4}}), \Gamma^{(0)} = \widetilde{\Theta}(\sigma_0\sigma_p d^{\frac{1}{2}}) = \widetilde{\Theta}(d^{-\frac{1}{4}+\epsilon})$, we have

$$\begin{aligned}
\eta\cdot T_r^* &\leq (1+\delta)\cdot\widetilde{\Theta}(d^{\frac{1}{4}}) + \eta(1+\delta)^2\Big(1+\frac{\log\widetilde{\Theta}(d^{\frac{1}{4}})}{\log(1+\delta)}\Big)\Big/(2p-1)\cdot\Theta(d)\\
&= \frac{3}{2(2p-1)}\widetilde{\Theta}(d^{-\frac{3}{4}}) + \frac{9}{4}\eta\Big(1+\frac{\widetilde{\Theta}(1)}{\log(\frac{3}{2})}\Big)\cdot\frac{\widetilde{\Theta}(d^{-1})}{2p-1} = \widetilde{\Theta}(d^{-\frac{3}{4}}),
\end{aligned} \tag{A.26}$$

and

$$\begin{aligned}
\eta\cdot T^* &\geq \frac{\widetilde{\Theta}(d^{\frac{1}{4}-\epsilon})}{1+\delta'} - \eta(1+\delta')^{-2}\Big(1+\frac{\log C}{\log(1+\delta')}\Big)\Big/\widetilde{\Theta}(d^{1-2\epsilon})\\
&= \frac{2}{3}\widetilde{\Theta}(d^{-\frac{3}{4}+\epsilon}) - \frac{4}{9}\eta\Big(1+\frac{\log C}{\log(\frac{3}{2})}\Big)\cdot\widetilde{\Theta}(d^{-1+\epsilon})\\
&= \widetilde{\Theta}(d^{-\frac{3}{4}+\epsilon}).
\end{aligned} \tag{A.27}$$

Therefore, combining (A.26) and (A.27), we have $\eta \cdot T^* \geq \widetilde{\Theta}(d^{-\frac{3}{4}+\epsilon}) > \widetilde{\Theta}(d^{-\frac{3}{4}}) \geq \eta \cdot T_r^*$, leading to $T^* > T_r^*$ for both $r \in \{-1. +1\}$. This indicates that when $\widehat{\Lambda}_1^{(t)}, \widehat{\Lambda}_{-1}^{(t)}$ reach $\Theta(1/m)$, $\Gamma^{(t)}$ remain the same magnitude as initialization. □

By leveraging tensor power method, we can also estimate the length of Stage I, i.e. T_1, T_{-1}, by applying tensor power method. To use tensor power method, we need to upper-bound the increasing speed of $\widehat{\Lambda}_r^{(t)}$. We have the following lemma:

Lemma A.15. For $r \in \{\pm 1\}$, we have with high probability that

$$\widehat{\Lambda}_r^{(t+1)} \geq (1-\eta\lambda) \cdot \widehat{\Lambda}_r^{(t)} + \eta \cdot 3\Big(p - \frac{1}{2} - o(1)\Big) \cdot \|\mathbf{v}\|_2^2 \cdot \big(\widehat{\Lambda}_r^{(t)}\big)_2^2,$$

$$\widehat{\Lambda}_r^{(t+1)} \leq (1-\eta\lambda) \cdot \widehat{\Lambda}_r^{(t)} + \eta \cdot 3\Big(p - \frac{1}{2} + o(1)\Big) \cdot \|\mathbf{v}\|_2^2 \cdot \big(\widehat{\Lambda}_r^{(t)}\big)_2^2.$$

Proof of lemma A.15. Let $j^* = \arg\max_{1\leq j\leq m} \langle \mathbf{w}_j^{(t+1)}, \mathbf{v}\rangle$ and note that $u_{j^*} = \mathbb{1}_{[1\leq j\leq m]} = \mathbb{1}_{[m+1\leq j\leq 2m]} = 1$, then we have

$$\begin{aligned}
\widehat{\Lambda}_1^{(t+1)} &= \langle \mathbf{w}_{j^*}^{(t+1)}, \mathbf{v}\rangle \\
&= (1-\eta\lambda) \cdot \langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v}\rangle + \frac{3\eta}{n_{\mathrm{l}}+n_{\mathrm{u}}} \underbrace{\Bigg(\sum_{i\in S_1} y_i \widehat{y}_i c_i^{(t)} [\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v}\rangle]_+^2 \|\mathbf{v}\|_2^2 + \sum_{i\in S_{-1}} y_i \widehat{y}_i c_i^{(t)} [-\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v}\rangle]_+^2 \|\mathbf{v}\|_2^2 \Bigg)}_{\clubsuit} \\
&\quad + \frac{3\eta}{n_{\mathrm{l}}+n_{\mathrm{u}}} \underbrace{\Bigg(\sum_{i\in S_1'} b_i^{(t)} [\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v}\rangle]_+^2 \|\mathbf{v}\|_2^2 + \sum_{i\in S_1'} b_i^{(t)} [-\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v}\rangle]_+^2 \|\mathbf{v}\|_2^2 \Bigg)}_{\bigstar}.
\end{aligned} \tag{A.28}$$

For ♣, according to lemma A.8, we have

$$\begin{aligned}
&\underbrace{\sum_{i\in S_1} y_i \widehat{y}_i c_i^{(t)} [\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v}\rangle]_+^2 \|\mathbf{v}\|_2^2 + \sum_{i\in S_{-1}} y_i \widehat{y}_i c_i^{(t)} [-\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v}\rangle]_+^2 \|\mathbf{v}\|_2^2}_{\clubsuit} \\
&= n_1 \cdot \Big(p - \frac{1}{2} \pm o(1)\Big) \cdot [\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v}\rangle]_+^2 \|\mathbf{v}\|_2^2 + n_{-1} \cdot \Big(p - \frac{1}{2} \pm o(1)\Big) \cdot [-\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v}\rangle]_+^2 \|\mathbf{v}\|_2^2 \\
&\leq n_1 \cdot \Big(p - \frac{1}{2} \pm o(1)\Big) \cdot \|\mathbf{v}\|_2^2 \cdot \big(\widehat{\Lambda}_1^{(t)}\big)^2 + n_{-1} \cdot \Big(p - \frac{1}{2} \pm o(1)\Big) \cdot \|\mathbf{v}\|_2^2 \cdot \big(\bar{\Lambda}_{-1}^{(t)}\big)^2 \\
&= n_1 \cdot \Big(p - \frac{1}{2} \pm o(1)\Big) \cdot \|\mathbf{v}\|_2^2 \cdot \big(\widehat{\Lambda}_1^{(t)}\big)^2,
\end{aligned} \tag{A.29}$$

where the last equality is due to $\widehat{\Lambda}_1^{(t)} = \omega(\bar{\Lambda}_{-1}^{(t)})$.

For ★, according to lemma A.7, we have

$$\begin{aligned}
&\underbrace{\sum_{i\in S_1'} b_i^{(t)} [\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v}\rangle]_+^2 \|\mathbf{v}\|_2^2 + \sum_{i\in S_1'} b_i^{(t)} [-\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v}\rangle]_+^2 \|\mathbf{v}\|_2^2}_{\bigstar} \\
&= n_1' \cdot \Big(\frac{1}{2} \pm o(1)\Big) \cdot [\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v}\rangle]_+^2 \|\mathbf{v}\|_2^2 + n_{-1}' \cdot \Big(\frac{1}{2} \pm o(1)\Big) \cdot [-\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v}\rangle]_+^2 \|\mathbf{v}\|_2^2 \\
&\leq n_1' \cdot \Big(\frac{1}{2} \pm o(1)\Big) \cdot \|\mathbf{v}\|_2^2 \cdot \big(\widehat{\Lambda}_1^{(t)}\big)^2 + n_{-1}' \cdot \Big(\frac{1}{2} \pm o(1)\Big) \cdot \|\mathbf{v}\|_2^2 \cdot \big(\bar{\Lambda}_{-1}^{(t)}\big)^2 \\
&= n_1' \cdot \Big(\frac{1}{2} \pm o(1)\Big) \cdot \|\mathbf{v}\|_2^2 \cdot \big(\widehat{\Lambda}_1^{(t)}\big)^2,
\end{aligned} \tag{A.30}$$

where the last equality is due to $\widehat{\Lambda}_1^{(t)} = \omega(\bar{\Lambda}_{-1}^{(t)})$.

Plugging (A.29) and (A.30) into (A.28), we have

$$\begin{aligned}
&\widehat{\Lambda}_1^{(t+1)} \\
&\leq (1-\eta\lambda)\cdot\widehat{\Lambda}_1^{(t)} + \frac{3\eta}{n_{\mathrm{l}}+n_{\mathrm{u}}}\bigg(n_1\cdot\Big(p-\frac{1}{2}\pm o(1)\Big)\cdot\|\mathbf{v}\|_2^2\cdot\big(\widehat{\Lambda}_1^{(t)}\big)^2 + n_1'\cdot\Big(\frac{1}{2}\pm o(1)\Big)\cdot\|\mathbf{v}\|_2^2\cdot\big(\widehat{\Lambda}_1^{(t)}\big)^2\bigg) \\
&= (1-\eta\lambda)\cdot\widehat{\Lambda}_1^{(t)} + \frac{3\eta n_1}{n_{\mathrm{l}}+n_{\mathrm{u}}}\cdot\Big(p-\frac{1}{2}\pm o(1)\Big)\cdot\|\mathbf{v}\|_2^2\cdot\big(\widehat{\Lambda}_1^{(t)}\big)^2 + \frac{3\eta n_1'}{n_{\mathrm{l}}+n_{\mathrm{u}}}\cdot\Big(\frac{1}{2}\pm o(1)\Big)\cdot\|\mathbf{v}\|_2^2\cdot\big(\widehat{\Lambda}_1^{(t)}\big)^2 \\
&= (1-\eta\lambda)\cdot\widehat{\Lambda}_1^{(t)} + 3\eta\cdot\bigg(\frac{n_1}{n_{\mathrm{l}}+n_{\mathrm{u}}}\cdot\Big(p-\frac{1}{2}\pm o(1)\Big) + \frac{n_1'}{n_{\mathrm{l}}+n_{\mathrm{u}}}\cdot\Big(\frac{1}{2}\pm o(1)\Big)\bigg)\cdot\|\mathbf{v}\|_2^2\cdot\big(\widehat{\Lambda}_1^{(t)}\big)^2 \\
&= (1-\eta\lambda)\cdot\widehat{\Lambda}_1^{(t)} + 3\eta\cdot\bigg(\underbrace{\frac{n_1}{n_{\mathrm{l}}+n_{\mathrm{u}}}\cdot\Big(p-\frac{1}{2}\Big) + \frac{n_1'}{n_{\mathrm{l}}+n_{\mathrm{u}}}\cdot\frac{1}{2}}_{\spadesuit}\pm o(1)\bigg)\cdot\|\mathbf{v}\|_2^2\cdot\big(\widehat{\Lambda}_1^{(t)}\big)^2. \quad \text{(A.31)}
\end{aligned}$$

Note that we have already proved in (A.9) that

$$\widehat{\Lambda}_1^{(t+1)} \leq (1-\eta\lambda)\cdot\widehat{\Lambda}_1^{(t)} + 3\eta\cdot\bigg(\underbrace{\frac{n_1}{n_{\mathrm{l}}+n_{\mathrm{u}}}\cdot\Big(p-\frac{1}{2}\Big) + \frac{n_1'}{n_{\mathrm{l}}+n_{\mathrm{u}}}\cdot\frac{1}{2}}_{\spadesuit}\pm o(1)\bigg)\cdot\|\mathbf{v}\|_2^2\cdot\big(\widehat{\Lambda}_1^{(t)}\big)^2. \quad \text{(A.32)}$$

Note we have already prove in (A.10) that

$$\underbrace{\frac{n_1}{n_{\mathrm{l}}+n_{\mathrm{u}}}\cdot\Big(p-\frac{1}{2}\Big) + \frac{n_1'}{n_{\mathrm{l}}+n_{\mathrm{u}}}\cdot\frac{1}{2}}_{\spadesuit} = \frac{1}{2}\cdot\Big(p-\frac{1}{2}\Big)\pm o(1)$$

Therefore, we have

$$\begin{aligned}
\widehat{\Lambda}_1^{(t+1)} &\geq (1-\eta\lambda)\cdot\widehat{\Lambda}_1^{(t)} + 3\eta\cdot\Big(p-\frac{1}{2}-o(1)\Big)\cdot\|\mathbf{v}\|_2^2\cdot\big(\widehat{\Lambda}_1^{(t)}\big)_2^2, \\
\widehat{\Lambda}_1^{(t+1)} &\leq (1-\eta\lambda)\cdot\widehat{\Lambda}_1^{(t)} + 3\eta\cdot\Big(p-\frac{1}{2}+o(1)\Big)\cdot\|\mathbf{v}\|_2^2\cdot\big(\widehat{\Lambda}_1^{(t)}\big)_2^2.
\end{aligned}$$

In a similar way, we can prove that

$$\begin{aligned}
\widehat{\Lambda}_{-1}^{(t+1)} &\geq (1-\eta\lambda)\cdot\widehat{\Lambda}_{-1}^{(t)} + 3\eta\cdot\Big(p-\frac{1}{2}-o(1)\Big)\cdot\|\mathbf{v}\|_2^2\cdot\big(\widehat{\Lambda}_{-1}^{(t)}\big)_2^2, \\
\widehat{\Lambda}_{-1}^{(t+1)} &\leq (1-\eta\lambda)\cdot\widehat{\Lambda}_{-1}^{(t)} + 3\eta\cdot\Big(p-\frac{1}{2}+o(1)\Big)\cdot\|\mathbf{v}\|_2^2\cdot\big(\widehat{\Lambda}_{-1}^{(t)}\big)_2^2,
\end{aligned}$$

which completes the proof of this lemma. □

By leveraging tensor power method given in lemma C.2, we have for $r\in\{\pm 1\}$ that

$$\begin{aligned}
\eta\cdot T_r^*\cdot C_t &= \sum_{t\geq 0,\widehat{\Lambda}_r^{(t)}\leq A}\eta\cdot C_t \leq \underbrace{\frac{1+\delta}{\widehat{\Lambda}_r^{(0)}}}_{(i)} + \underbrace{\eta(1+\delta)^2\Big(1+\frac{\log\big(A/\widehat{\Lambda}_r^{(0)}\big)}{\log(1+\delta)}\Big)}_{(ii)}, \\
\eta\cdot T_r^*\cdot C_t' &= \sum_{t\geq 0,\widehat{\Lambda}_r^{(t)}\leq A}\eta\cdot C_t' \geq \underbrace{\frac{1-\frac{\widehat{\Lambda}_r^{(0)}}{A}}{(1+\delta)\cdot\widehat{\Lambda}_r^{(0)}}}_{(iii)} - \underbrace{\eta(1+\delta)^{-2}\Big(1+\frac{\log\big(A/\widehat{\Lambda}_r^{(0)}\big)}{\log(1+\delta)}\Big)}_{(iv)},
\end{aligned}$$

where C_t is taken as $3\big(p-\frac{1}{2}-o(1)\big)\cdot\|\mathbf{v}\|_2^2/m$ and C_t' is taken as $3\big(p-\frac{1}{2}+o(1)\big)\cdot\|\mathbf{v}\|_2^2/m$ according to lemma A.15. Taking $\delta=\frac{1}{k}$, $A=\Theta(1/m)$ and note that terms $(ii),(iv)$ are respectively dominated by terms $(i),(iii)$, we have

$$\frac{k+1}{k\cdot\widehat{\Lambda}_r^{(0)}} = \frac{1+\delta}{\widehat{\Lambda}_r^{(0)}} \geq \eta\cdot T_r^*\cdot C_t \geq \frac{1}{(1+\delta)\cdot\widehat{\Lambda}_r^{(0)}} = \frac{k}{(k+1)\cdot\widehat{\Lambda}_r^{(0)}}. \quad \text{(A.33)}$$

Letting $k\to\infty$, we have $\eta\cdot T_r^* = (\widehat{\Lambda}_r^{(0)}\cdot C_t)^{-1}\pm\{\text{lower order terms}\} = \widetilde{\Theta}(d^{-\frac{3}{4}})$ for $r\in\{+1,-1\}$.

Lemma A.16 (Training time of pretext task)**.** For $r \in \{\pm 1\}$, let T_r be the first iteration that $\widehat{\Lambda}_r^{(t)}$ reaches $\Theta(1/m)$ respectively. Then $T_r = \widetilde{\Theta}(\sigma_0^{-1}\|\mathbf{v}\|_2^{-3})/\eta = \widetilde{\Theta}(d^{-\frac{3}{4}})/\eta$ for all $r \in \{\pm 1\}$.

The discussion above verifies lemma A.16 and the latter part of lemma 5.3 and provides a clear understanding about how $\widehat{\Lambda}_r^{(t)}, \bar{\Lambda}_r^{(t)}$ varies within the iteration range $[0, T_r]$ for $r \in \{+1, -1\}$. Note that the iteration numbers when $\widehat{\Lambda}_1^{(t)}$ and $\widehat{\Lambda}_{-1}^{(t)}$ reaches $\Theta(1/m)$ (T_1 and T_{-1}) are different.

However, since T_{-1} and T_1 have the same magnitude, it remains clear that although $T_1 \neq T_{-1}$ (wlog, assume $T_1 < T_{-1}$), we still have $\widehat{\Lambda}_1^{(t)} = \widetilde{\Theta}(1)$ and $\bar{\Lambda}_1^{(t)} = o(d^{-\frac{1}{4}})$ within the iteration range $[T_1, T_{-1}]$, since Stage II of GD (off-diagonal feature learning) also costs time no less than order $\widetilde{\Theta}(d^{-\frac{3}{4}}/\eta)$, which is higher order than $|T_1 - T_{-1}|$. Therefore, at time $T_0 := \max\{T_1, T_{-1}\}$, off-diagonal $\bar{\Lambda}_1^{(t)}, \bar{\Lambda}_{-1}^{(t)}$ still remain initialization magnitude $\widetilde{O}(d^{-\frac{1}{4}})$, $\Gamma_1^{(t)}, \Gamma_{-1}^{(t)}$ remain initialization magnitude $\widetilde{O}(d^{-\frac{1}{4}+\epsilon})$, while on-diagonal $\widehat{\Lambda}_1^{(t)}, \widehat{\Lambda}_{-1}^{(t)}$ reach and then remain $\widetilde{\Theta}(1)$.

A.5 Proof of lemma 5.2

If we only use labeled data S' for the optimization of CNN, according to lemma B.1, we have

$$\begin{aligned}\mathbf{w}_j^{(t+1)} &= \mathbf{w}_j^{(t)} - \nabla_{\mathbf{w}_j} L_{S'}(\mathbf{W}) \\ &= (1-\eta\lambda)\cdot \mathbf{w}_j^{(t)} + \frac{3\eta u_j}{n_1}\sum_{i=1}^{n_1} b_i^{(t)} y_i' \big([\langle \mathbf{w}_j, y_i'\cdot \mathbf{v}\rangle]_+^2 \cdot y_i'\cdot \mathbf{v} + [\langle \mathbf{w}_j, \boldsymbol{\xi}_i'\rangle]_+^2 \cdot \boldsymbol{\xi}_i'\big),\end{aligned}$$

where $u_j := \mathbb{1}_{[1\le j\le m]} - \mathbb{1}_{[m+1\le 2m]}$, $b_i^{(t)} = -\ell'(y_i'\cdot f_{\mathbf{W}}(\mathbf{x}_i')) = \exp[-y_i'\cdot f_{\mathbf{W}}(\mathbf{x}_i')]/(1+\exp[-y_i'\cdot f_{\mathbf{W}}(\mathbf{x}_i')])$.

Notice that $\mathbf{v}$ and $\boldsymbol{\xi}_i'$ are orthogonal to each other, we have

$$\begin{aligned}\langle \mathbf{w}_j^{(t+1)}, \mathbf{v}\rangle &= (1-\eta\lambda)\cdot\langle \mathbf{w}_j^{(t)}, \mathbf{v}\rangle + \frac{3\eta u_j}{n_1}\sum_{i=1}^{n_1} b_i^{(t)}\cdot[\langle \mathbf{w}_j^{(t)}, y_i'\cdot\mathbf{v}\rangle]_+^2\cdot\|\mathbf{v}\|_2^2, \\ \langle \mathbf{w}_j^{(t+1)}, \boldsymbol{\xi}_l'\rangle &= (1-\eta\lambda)\cdot\langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_i\rangle + \frac{3\eta u_j}{n_1}\sum_{i=1}^{n_1} b_i^{(t)} y_i'\cdot[\langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_i'\rangle]_+^2\cdot\langle \boldsymbol{\xi}_i', \boldsymbol{\xi}_l'\rangle, i\in[n_1].\end{aligned}$$

Let T_i' be the first iteration that $\Gamma_i'^{(t)}$ reaches $\Theta(1/m)$, then we have following lemma:

Lemma A.17. As long as $\Gamma_i'^{(t)} \le \Theta(1/m)$, $b_i^{(t)} := -\ell'(y_i'\cdot f_{\mathbf{W}^{(t)}}(\mathbf{x}_i'))$ will remain $1/2 \pm o(1)$.

Proof of lemma A.17. Note that $\ell(z) = \log(1+\exp(-z))$ and $-\ell'(z) = \exp(-z)/\big(1+\exp(-z)\big)$, and without loss of generality assuming $y_i' = 1$, we can express $b_i^{(t)}$ as follow:

$$b_i^{(t)} = -\ell'(f_{\mathbf{W}^{(t)}}(\mathbf{x}_i')) = \frac{e^{\sum_{j=m+1}^{2m}[\sigma(\langle \mathbf{w}_j^{(t)},\mathbf{v}\rangle)+\sigma(\langle \mathbf{w}_j^{(t)},\boldsymbol{\xi}_i'\rangle)]}}{e^{\sum_{j=1}^{m}[\sigma(\langle \mathbf{w}_j^{(t)},\mathbf{v}\rangle)+\sigma(\langle \mathbf{w}_j^{(t)},\boldsymbol{\xi}_i'\rangle)]} + e^{\sum_{j=m+1}^{2m}[\sigma(\langle \mathbf{w}_j^{(t)},\mathbf{v}\rangle)+\sigma(\langle \mathbf{w}_j^{(t)},\boldsymbol{\xi}_i'\rangle)]}},$$

Since $\sigma(\langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}\rangle)$ will dominate $\sigma(\langle \mathbf{w}_j^{(t)}, \mathbf{v}\rangle)$, which will be proved later by using *tensor power method*, we have

$$b_i^{(t)} = -\ell'(f_{\mathbf{W}^{(t)}}(\mathbf{x}_i')) = \frac{e^{\sum_{j=m+1}^{2m}[\sigma(\langle \mathbf{w}_j^{(t)},\mathbf{v}\rangle)+\sigma(\langle \mathbf{w}_j^{(t)},\boldsymbol{\xi}_i'\rangle)]}}{e^{\sum_{j=1}^{m}\sigma(\langle \mathbf{w}_j^{(t)},\boldsymbol{\xi}_i'\rangle)\{+\text{lower order term}\}} + e^{\sum_{j=m+1}^{2m}[\sigma(\langle \mathbf{w}_j^{(t)},\mathbf{v}\rangle)+\sigma(\langle \mathbf{w}_j^{(t)},\boldsymbol{\xi}_i'\rangle)]}},$$

On the one side,

$$\begin{aligned}b_i^{(t)} &\ge \frac{1}{e^{\sum_{j=1}^{m}\sigma(\langle \mathbf{w}_j^{(t)},\boldsymbol{\xi}_i'\rangle)\{+\text{lower order term}\}}+1} \\ &\ge \frac{1}{e^{m(\Gamma_i'^{(t)})^3\{+\text{lower order term}\}}+1} \\ &\ge \frac{1}{e^{\Theta(m^{-2})}+1} = \frac{1}{2+o(1)} = \frac{1}{2} - o(1).\end{aligned}$$

On the other side, according to lemma 5.4, we have $\bar{\Lambda}_1^{(t)} = \widetilde{O}(d^{-\frac{1}{4}})$, it follows that

$$
\begin{aligned}
b_i^{(t)} &\leq \frac{e^{m(\bar{\Lambda}_1^{(t)})^3+o(1)}}{e^{\sum_{j=1}^m \sigma(\langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_i' \rangle) + \{\text{lower order term}\}} + e^{m(\bar{\Lambda}_1^{(t)})^3+o(1)}} \\
&= \frac{1+o(1)}{e^{\sum_{j=1}^m \sigma(\langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_i' \rangle) + \{\text{lower order term}\}} + 1 + o(1)} \\
&\leq \frac{1+o(1)}{1+1+o(1)} = \frac{1}{2} + o(1).
\end{aligned}
$$

Therefore, we have $b_i^{(t)} = 1/2 \pm o(1)$ and the other case of $y_i = -1$ can be proved in a similar way. □

With the help of above lemma, we are now ready to prove lemma 5.2.

Proof of lemma 5.2. Let $j^* = \arg\max_{1 \leq j \leq m} \langle \mathbf{w}_j^{(t+1)}, \mathbf{v} \rangle$ and note that $u_j = 1$, according to lemma A.17, we have

$$
\begin{aligned}
\widehat{\Lambda}_1^{(t+1)} &= \langle \mathbf{w}_{j^*}^{(t+1)}, \mathbf{v} \rangle \\
&= (1-\eta\lambda) \cdot \langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v} \rangle + \frac{3\eta}{n_1} \sum_{i=1}^{n_1} b_i^{(t)} [\langle \mathbf{w}_{j^*}^{(t)}, y_i' \cdot \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 \\
&= (1-\eta\lambda) \cdot \langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v} \rangle + \frac{3\eta}{n_1} \sum_{i \in S_1'} b_i^{(t)} [\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 + \frac{3\eta}{n_1} \sum_{i \in S_{-1}'} b_i^{(t)} [\langle \mathbf{w}_{j^*}^{(t)}, -\mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 \\
&= (1-\eta\lambda) \cdot \langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v} \rangle + \frac{3\eta}{n_1} \underbrace{\sum_{i \in S_1'} \Big(\frac{1}{2} \pm o(1)\Big) [\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2}_{\clubsuit} \\
&\quad + \frac{3\eta}{n_1} \underbrace{\sum_{i \in S_{-1}'} \Big(\frac{1}{2} \pm o(1)\Big) [\langle \mathbf{w}_{j^*}^{(t)}, -\mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2}_{\bigstar}
\end{aligned}
\tag{A.34}
$$

For ♣, we have

$$
\begin{aligned}
\underbrace{\sum_{i \in S_1'} \Big(\frac{1}{2} \pm o(1)\Big) [\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2}_{\clubsuit} &= n_1' \cdot \Big(\frac{1}{2} \pm o(1)\Big) \cdot [\langle \mathbf{w}_{j^*}^{(t)}, \mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 \\
&\leq n_1' \cdot \Big(\frac{1}{2} \pm o(1)\Big) \cdot \|\mathbf{v}\|_2^2 \cdot (\widehat{\Lambda}_1^{(t)})^2.
\end{aligned}
\tag{A.35}
$$

For ★, we have

$$
\begin{aligned}
\underbrace{\sum_{i \in S_{-1}'} \Big(\frac{1}{2} \pm o(1)\Big) [\langle \mathbf{w}_{j^*}^{(t)}, -\mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2}_{\bigstar} &= n_{-1}' \cdot \Big(\frac{1}{2} \pm o(1)\Big) \cdot [\langle \mathbf{w}_{j^*}^{(t)}, -\mathbf{v} \rangle]_+^2 \|\mathbf{v}\|_2^2 \\
&\leq n_{-1}' \cdot \Big(\frac{1}{2} \pm o(1)\Big) \cdot \|\mathbf{v}\|_2^2 \cdot (\bar{\Lambda}_{-1}^{(t)})^2.
\end{aligned}
\tag{A.36}
$$

By plugging (A.35) and (A.36) in (A.34), and according to lemma A.10, we have with probability at least $1-4\delta$ that

$$\begin{aligned}
\widehat{\Lambda}_1^{(t+1)} &\leq (1-\eta\lambda)\cdot\widehat{\Lambda}_1^{(t)} + \frac{3\eta}{n_1}\bigg(n_1'\cdot\Big(\frac{1}{2}\pm o(1)\Big)\cdot\|\mathbf{v}\|_2^2\cdot(\widehat{\Lambda}_1^{(t)})^2 + n_{-1}'\cdot\Big(\frac{1}{2}\pm o(1)\Big)\cdot\|\mathbf{v}\|_2^2\cdot(\bar{\Lambda}_{-1}^{(t)})^2\bigg) \\
&\leq (1-\eta\lambda)\cdot\widehat{\Lambda}_1^{(t)} + \frac{3\eta}{n_1}\bigg(\Big(\frac{n_1}{2}+\sqrt{\frac{n_1}{2}\log\frac{1}{\delta}}\Big)\cdot\Big(\frac{1}{2}\pm o(1)\Big)\cdot\|\mathbf{v}\|_2^2\cdot(\widehat{\Lambda}_1^{(t)})^2 \\
&\qquad\qquad + \Big(\frac{n_1}{2}+\sqrt{\frac{n_1}{2}\log\frac{1}{\delta}}\Big)\cdot\Big(\frac{1}{2}\pm o(1)\Big)\cdot\|\mathbf{v}\|_2^2\cdot(\bar{\Lambda}_{-1}^{(t)})^2\bigg) \\
&= (1-\eta\lambda)\cdot\widehat{\Lambda}_1^{(t)} + 3\eta\bigg(\Big(\frac{1}{4}\pm o(1)\Big)\cdot\|\mathbf{v}\|_2^2\cdot(\widehat{\Lambda}_1^{(t)})^2 + \Big(\frac{1}{4}\pm o(1)\Big)\cdot\|\mathbf{v}\|_2^2\cdot(\bar{\Lambda}_{-1}^{(t)})^2\bigg) \\
&= (1-\eta\lambda)\cdot\widehat{\Lambda}_1^{(t)} + \eta\cdot\Theta(d)\cdot\Big((\widehat{\Lambda}_1^{(t)})^2 + (\bar{\Lambda}_{-1}^{(t)})^2\Big).
\end{aligned}$$

And we can prove in the same way that with probability at least $1-4\delta$ we have

$$\widehat{\Lambda}_{-1}^{(t+1)} \leq (1-\eta\lambda)\cdot\widehat{\Lambda}_{-1}^{(t)} + \eta\cdot\Theta(d)\cdot\Big((\widehat{\Lambda}_{-1}^{(t)})^2 + (\bar{\Lambda}_1^{(t)})^2\Big).$$

Let $j^\star = \arg\max_{m+1\leq j\leq 2m}\langle\mathbf{w}_j^{(t+1)},\mathbf{v}\rangle$ and note that $u_j=-1$, we have

$$\begin{aligned}
\bar{\Lambda}_1^{(t+1)} &= \langle\mathbf{w}_{j^\star}^{(t+1)},\mathbf{v}\rangle \\
&= (1-\eta\lambda)\cdot\langle\mathbf{w}_{j^\star}^{(t)},\mathbf{v}\rangle - \frac{3\eta}{n_1}\sum_{i=1}^{n_1} b_i^{(t)}[\langle\mathbf{w}_{j^\star}^{(t)}, y_i'\cdot\mathbf{v}\rangle]_+^2\|\mathbf{v}\|_2^2 \\
&\leq (1-\eta\lambda)\cdot\langle\mathbf{w}_{j^\star}^{(t)},\mathbf{v}\rangle \\
&\leq (1-\eta\lambda)\cdot\bar{\Lambda}_1^{(t)}.
\end{aligned} \tag{A.37}$$

And we can prove in the same way that $\bar{\Lambda}_{-1}^{(t+1)} \leq (1-\eta\lambda)\cdot\bar{\Lambda}_{-1}^{(t)}$.

Next, we consider the increasing rate of $\Gamma_l'^{(t)}$ where $l\in[n_1]$ is fixed. If $y_l=1$, let $j^\natural = \mathrm{argmax}_{1\leq j\leq m}\langle\mathbf{w}_j^{(t)},\boldsymbol{\xi}_l'\rangle$ and note that $u_j=1$, we have

$$\begin{aligned}
\Gamma_l'^{(t+1)} &\geq \langle\mathbf{w}_{j^\natural}^{(t+1)},\boldsymbol{\xi}_l'\rangle \\
&= (1-\eta\lambda)\cdot\langle\mathbf{w}_{j^\natural}^{(t)},\boldsymbol{\xi}_l'\rangle + \frac{3\eta}{n_1}\sum_{i=1}^{n_1} b_i^{(t)} y_i'\cdot[\langle\mathbf{w}_{j^\natural}^{(t)},\boldsymbol{\xi}_i'\rangle]_+^2\cdot\langle\boldsymbol{\xi}_i',\boldsymbol{\xi}_l'\rangle \\
&= (1-\eta\lambda)\cdot\langle\mathbf{w}_{j^\natural}^{(t)},\boldsymbol{\xi}_l'\rangle + \frac{3\eta}{n_1}b_l^{(t)}[\langle\mathbf{w}_{j^\natural}^{(t)},\boldsymbol{\xi}_l'\rangle]_+^2\|\boldsymbol{\xi}_l'\|_2^2 + \frac{3\eta}{n_1}\sum_{i\in[n_1],i\neq l} b_i^{(t)} y_i'[\langle\mathbf{w}_{j^\natural}^{(t)},\boldsymbol{\xi}_i'\rangle]_+^2\langle\boldsymbol{\xi}_i',\boldsymbol{\xi}_l'\rangle \\
&= (1-\eta\lambda)\cdot\langle\mathbf{w}_{j^\natural}^{(t)},\boldsymbol{\xi}_l'\rangle + \frac{3\eta}{n_1}b_l^{(t)}[\langle\mathbf{w}_{j^\natural}^{(t)},\boldsymbol{\xi}_l'\rangle]_+^2\|\boldsymbol{\xi}_l'\|_2^2\{\pm\text{ lower order terms}\} \\
&\geq (1-\eta\lambda)\cdot\Gamma_l'^{(t)} + \frac{3\eta}{n_1}\cdot\Big(\frac{1}{2}-o(1)\Big)\cdot\|\boldsymbol{\xi}_l'\|_2^2\cdot(\Gamma_l'^{(t)})^2 \\
&= (1-\eta\lambda)\cdot\Gamma_l'^{(t)} + \eta\cdot\widetilde{\Theta}(d^{1+2\epsilon})\cdot(\Gamma_l'^{(t)})^2.
\end{aligned} \tag{A.38}$$

If $y_l = -1$, let $j^\sharp = \text{argmax}_{m+1\le j\le 2m}\langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_l'\rangle$ and note that $u_j = -1$, we have

$$\begin{aligned}
\Gamma_l'^{(t+1)} &\ge \langle \mathbf{w}_{j^\sharp}^{(t+1)}, \boldsymbol{\xi}_l'\rangle \\
&= (1-\eta\lambda)\cdot\langle \mathbf{w}_{j^\sharp}^{(t)}, \boldsymbol{\xi}_l'\rangle - \frac{3\eta}{n_1}\sum_{i=1}^{n_1} b_i^{(t)} y_i' \cdot [\langle \mathbf{w}_{j^\sharp}^{(t)}, \boldsymbol{\xi}_i'\rangle]_+^2 \cdot \langle \boldsymbol{\xi}_i', \boldsymbol{\xi}_l'\rangle \\
&= (1-\eta\lambda)\cdot\langle \mathbf{w}_{j^\sharp}^{(t)}, \boldsymbol{\xi}_l'\rangle + \frac{3\eta}{n_1} b_l^{(t)}[\langle \mathbf{w}_{j^\sharp}^{(t)}, \boldsymbol{\xi}_l'\rangle]_+^2 \|\boldsymbol{\xi}_l'\|_2^2 - \frac{3\eta}{n_1}\sum_{i\in[n_1], i\ne l} b_i^{(t)} y_i' [\langle \mathbf{w}_{j^\sharp}^{(t)}, \boldsymbol{\xi}_i'\rangle]_+^2 \langle \boldsymbol{\xi}_i', \boldsymbol{\xi}_l'\rangle \\
&= (1-\eta\lambda)\cdot\langle \mathbf{w}_{j^\sharp}^{(t)}, \boldsymbol{\xi}_l'\rangle + \frac{3\eta}{n_1} b_l^{(t)}[\langle \mathbf{w}_{j^\sharp}^{(t)}, \boldsymbol{\xi}_l'\rangle]_+^2 \|\boldsymbol{\xi}_l'\|_2^2 \{\pm \text{ lower order terms}\} \\
&\ge (1-\eta\lambda)\cdot\Gamma_l'^{(t)} + \frac{3\eta}{n_1}\cdot\Big(\frac{1}{2} - o(1)\Big)\cdot\|\boldsymbol{\xi}_l'\|_2^2\cdot(\Gamma_l'^{(t)})^2 \\
&= (1-\eta\lambda)\cdot\Gamma_l'^{(t)} + \eta\cdot\widetilde{\Theta}(d^{1+2\epsilon})\cdot(\Gamma_l'^{(t)})^2.
\end{aligned} \tag{A.39}$$

According to (A.38) and (A.39), we always have

$$\Gamma_l'^{(t+1)} \ge (1-\eta\lambda)\cdot\Gamma_l'^{(t)} + \eta\cdot\widetilde{\Theta}(d^{1+2\epsilon})\cdot(\Gamma_l'^{(t)})^2.$$

□

A.6 Proof of Lemma 5.4

By applying Lemma C.2 to $\Gamma_i^{(t)}$ and taking $C_t = \widetilde{\Theta}(d^{1+2\epsilon})$, $\delta = 1/2$, $A = \Theta(1/m)$, we have

$$\sum_{t\ge 0, \Gamma_i^{(t)}\le A} \eta\cdot C_t \le \frac{2}{3\Gamma_i^{(0)}} + \eta\cdot\frac{9}{4}\cdot\Big(1 + \frac{\log(A/\Gamma_i^{(0)})}{\log(3/2)}\Big) = \widetilde{\Theta}(d^{\frac{1}{4}-\epsilon}).$$

And note the definition of T_i', we have

$$\eta\cdot T_i' = \frac{\widetilde{\Theta}(d^{\frac{1}{4}-\epsilon})}{C_t} = \widetilde{\Theta}(d^{-\frac{3}{4}-3\epsilon}). \tag{A.40}$$

In Lemma 5.2, we have already prove that

$$\begin{aligned}
&\widehat{\Lambda}_r^{(t+1)} \le (1-\eta\lambda)\cdot\widehat{\Lambda}_r^{(t)} + \eta\cdot\Theta(d)\cdot\Big((\widehat{\Lambda}_r^{(t)})^2 + (\bar{\Lambda}_{-r}^{(t)})^2\Big), \\
&\widehat{\Lambda}_r^{(t+1)} \le (1-\eta\lambda)\cdot\widehat{\Lambda}_r^{(t+1)}, r\in\{\pm 1\}.
\end{aligned} \tag{A.41}$$

Define $\Lambda^{(t)} := \max_{r\in\{\pm1\}}\{\widehat{\Lambda}_r^{(t)}, \bar{\Lambda}_r^{(t)}\}$, according to (A.41), we have

$$\Lambda^{(t+1)} \le (1-\eta\lambda)\cdot\Lambda^{(t)} + \eta\cdot\Theta(d)\cdot(\Lambda^{(t)})^2.$$

By applying Lemma C.2 to $\Lambda^{(t)}$, and taking $C_t = \Theta(d)$, $\delta = 1/2$, $A = C\cdot\Lambda^{(0)}$, where A is a large constant, we have

$$\sum_{t\ge 0, \Lambda^{(t)}\le A} \eta\cdot C_t \ge \frac{1 - \frac{\Lambda^{(0)}}{A}}{\frac{3\Lambda^{(0)}}{2}} - \eta\cdot\frac{4}{9}\cdot\Big(1 + \frac{\log(A/\Lambda^{(0)})}{\log(3/2)}\Big) = \widetilde{\Theta}(d^{\frac{1}{4}}).$$

Let T' be the first iteration that $\Lambda^{(t)}$ reaches $C\cdot\Lambda^{(0)}$, then we have

$$\eta\cdot T' = \frac{\widetilde{\Theta}(d^{\frac{1}{4}})}{C_t} = \widetilde{\Theta}(d^{-\frac{3}{4}}). \tag{A.42}$$

According to (A.40) and (A.42), we have $T' = \omega(T_i')$, which indicates that when $\Gamma_i^{(t)}$ reaches $\Theta(1/m)$, $\Lambda^{(t)}$ remains initialization magnitude $\widetilde{\Theta}(d^{-\frac{1}{4}})$.

A.7 EMPIRICAL, TEST ERROR AND LOSS FOR EARLY STOPPED CLASSIFIER

Assume the accuracy of pseudolabeler p is larger than $1/2$. We first estimate the empirical loss for early stopped classifier $f_{\mathbf{W}^{(T_0)}}$, where $T_0 = \max_{r\in\{\pm 1\}}\{T_r\}$ and T_r is defined as the first iteration that $\widehat{\Lambda}_r^{(t)}$ reaches $\Theta(1/m)$. According to section A.4.3 and lemma A.14, we have $\widehat{\Lambda}_r^{(T_0)} = \widetilde{\Theta}(1), \bar{\Lambda}_r^{(T_0)} = \widetilde{O}(d^{-\frac{1}{4}}), \Gamma^{(t)} = \widetilde{O}(d^{-\frac{1}{4}+\epsilon})$, for $r \in \{+1, -1\}$. We have the following lemma:

Lemma A.18. Early stopped classifier $f_{\mathbf{W}^{(T_0)}}(\mathbf{x})$ possesses following properties:

1. Training error of early stopped classifier $f_{\mathbf{W}^{(T_0)}}(\mathbf{x})$ is asymptotically $1-p$: $\frac{1}{n_{\mathrm{u}}+n_{\mathrm{l}}}\big(\sum_{i=1}^{n_{\mathrm{u}}} \mathbb{1}[\widehat{y}_i \cdot f_{\mathbf{W}^{(T_0)}}(\mathbf{x}_i) \le 0] + \sum_{i=1}^{n_{\mathrm{l}}} \mathbb{1}[y_i' \cdot f_{\mathbf{W}^{(T_0)}}(\mathbf{x}_i') \le 0]\big) = 1 - p \pm o(1)$.
2. Test error is nearly $1-p$, if we use pseudo label $\widehat{y}$ generated by pseudolabeler as target: $\mathbb{P}_{(\mathbf{x},y)\sim\mathcal{D},\widehat{y}\sim y\cdot\mathcal{B}(p)}[\widehat{y}\cdot f_{\mathbf{W}^{(T_0)}}(\mathbf{x}) \le 0] = 1 - p \pm o(1)$.
3. Test error is nearly 0, if we use true label y as target: $\mathbb{P}_{(\mathbf{x},y)\sim\mathcal{D}}[y \cdot f_{\mathbf{W}^{(T_0)}}(\mathbf{x}) \le 0] = o(1)$ and hence $\operatorname{sign} f_{\mathbf{W}^{(T_0)}}(\mathbf{x}) = \operatorname{sign}(y)$ with high probability,

where p is the accuracy of the pseudolabeler. We can regard p as the probability that $\mathbf{x}_i$ is paired with true label y_i, $1-p$ is the probability that $\mathbf{x}_i$ is paired with wrong label $-y_i$.

Proof of lemma A.18. Recall the definition of $f_{\mathbf{W}}$ in (3.1) that

$$
\begin{aligned}
f_{\mathbf{W}^{(T_0)}}(\mathbf{x}_i) = &\sum_{j=1}^{m}\Big[\sigma\big(\langle \mathbf{w}_j^{(T_0)}, y_i\cdot\mathbf{v}\rangle\big) + \sigma\big(\langle \mathbf{w}_j^{(T_0)}, \boldsymbol{\xi}_i\rangle\big)\Big] \\
&- \sum_{j=m+1}^{2m}\Big[\sigma\big(\langle \mathbf{w}_j^{(T_0)}, y_i\cdot\mathbf{v}\rangle\big) + \sigma\big(\langle \mathbf{w}_j^{(T_0)}, \boldsymbol{\xi}_i\rangle\big)\Big].
\end{aligned}
$$

According to section A.4.3 and lemma A.14, we have $\widehat{\Lambda}_r^{(T_0)} = \widetilde{\Theta}(1), \bar{\Lambda}_r^{(T_0)} = \widetilde{O}(d^{-\frac{1}{4}}), \Gamma^{(t)} = \max\big\{\max_{i\in[n_{\mathrm{u}}]}\Gamma_i^{(t)}, \max_{i\in[n_{\mathrm{l}}]}\Gamma_i'^{(t)}\big\} = \widetilde{O}(d^{-\frac{1}{4}+\epsilon})$, for $r \in \{+1,-1\}$. If $y_i = 1$, we have following lower bound for $f_{\mathbf{W}^{(T_0)}}(\mathbf{x}_i)$

$$
\begin{aligned}
f_{\mathbf{W}^{(T_0)}}(\mathbf{x}_i) &= \sum_{j=1}^{m}\Big[\sigma\big(\langle \mathbf{w}_j^{(T_0)}, \mathbf{v}\rangle\big) + \sigma\big(\langle \mathbf{w}_j^{(T_0)}, \boldsymbol{\xi}_i\rangle\big)\Big] - \sum_{j=m+1}^{2m}\Big[\sigma\big(\langle \mathbf{w}_j^{(T_0)}, \mathbf{v}\rangle\big) + \sigma\big(\langle \mathbf{w}_j^{(T_0)}, \boldsymbol{\xi}_i\rangle\big)\Big] \\
&\ge \big(\widehat{\Lambda}_1^{(T_0)}\big)^3 + \big(\Gamma_i^{(T_0)}\big)^3 - m\big(\bar{\Lambda}_1^{(T_0)}\big)^3 - m\big(\Gamma_i^{(T_0)}\big)^3 \\
&\ge \big(\widehat{\Lambda}_1^{(T_0)}\big)^3\{-\text{ lower order terms}\},
\end{aligned}
$$

and following upper bound for $f_{\mathbf{W}^{(T_0)}}(\mathbf{x}_i)$:

$$
\begin{aligned}
f_{\mathbf{W}^{(T_0)}}(\mathbf{x}_i) &= \sum_{j=1}^{m}\Big[\sigma\big(\langle \mathbf{w}_j^{(T_0)}, \mathbf{v}\rangle\big) + \sigma\big(\langle \mathbf{w}_j^{(T_0)}, \boldsymbol{\xi}_i\rangle\big)\Big] - \sum_{j=m+1}^{2m}\Big[\sigma\big(\langle \mathbf{w}_j^{(T_0)}, \mathbf{v}\rangle\big) + \sigma\big(\langle \mathbf{w}_j^{(T_0)}, \boldsymbol{\xi}_i\rangle\big)\Big] \\
&\le m(\widehat{\Lambda}_1^{(T_0)})^3 + m(\Gamma_i^{(T_0)})^3 - \big(\bar{\Lambda}_1^{(T_0)}\big)^3 - \big(\Gamma_i^{(T_0)}\big)^3 \\
&\le (\widehat{\Lambda}_1^{(T_0)})^3\{+\text{ lower order terms}\}.
\end{aligned}
$$

If $y_i = -1$, we have following upper bound for $f_{\mathbf{W}^{(T_0)}}(\mathbf{x}_i)$:

$$
\begin{aligned}
f_{\mathbf{W}^{(T_0)}}(\mathbf{x}_i) &= \sum_{j=1}^{m}\Big[\sigma\big(-\langle \mathbf{w}_j^{(T_0)}, \mathbf{v}\rangle\big) + \sigma\big(\langle \mathbf{w}_j^{(T_0)}, \boldsymbol{\xi}_i\rangle\big)\Big] - \sum_{j=m+1}^{2m}\Big[\sigma\big(-\langle \mathbf{w}_j^{(T_0)}, \mathbf{v}\rangle\big) + \sigma\big(\langle \mathbf{w}_j^{(T_0)}, \boldsymbol{\xi}_i\rangle\big)\Big] \\
&\le m\big(\bar{\Lambda}_{-1}^{(T_0)}\big)^3 + m\big(\Gamma_i^{(T_0)}\big)^3 - \big(\widehat{\Lambda}_{-1}^{(T_0)}\big)^3 - \big(\Gamma_i^{(T_0)}\big)^3 \\
&\le -\big(\widehat{\Lambda}_{-1}^{(T_0)}\big)^3\{+\text{ lower order terms}\},
\end{aligned}
$$

and following lower bound for $f_{\mathbf{W}^{(T_0)}}(\mathbf{x}_i)$:

$$\begin{aligned} f_{\mathbf{W}^{(T_0)}}(\mathbf{x}_i) &= \sum_{j=1}^{m}\Big[\sigma\big(-\langle \mathbf{w}_j^{(T_0)}, \mathbf{v}\rangle\big) + \sigma\big(\langle \mathbf{w}_j^{(T_0)}, \boldsymbol{\xi}_i\rangle\big)\Big] - \sum_{j=m+1}^{2m}\Big[\sigma\big(-\langle \mathbf{w}_j^{(T_0)}, \mathbf{v}\rangle\big) + \sigma\big(\langle \mathbf{w}_j^{(T_0)}, \boldsymbol{\xi}_i\rangle\big)\Big] \\ &\geq \big(\bar{\Lambda}_{-1}^{(T_0)}\big)^3 + \big(\Gamma_i^{(T_0)}\big)^3 - m\big(\widehat{\Lambda}_{-1}^{(T_0)}\big)^3 - m\big(\Gamma_i^{(T_0)}\big)^3 \\ &\geq -m\big(\bar{\Lambda}_{-1}^{(T_0)}\big)^3\{-\text{ lower order terms}\}. \end{aligned}$$

Therefore, for unlabeled data, we have $y_i \cdot f_{\mathbf{W}^{(T_0)}}(\mathbf{x}_i) \in \big[\big(1-o(1)\big)\cdot(\widehat{\Lambda}_{y_i}^{(T_0)})^3, \big(m+o(1)\big)\cdot(\widehat{\Lambda}_{y_i}^{(T_0)})^3\big]$ and hence $\operatorname{sign}\big(f_{\mathbf{W}^{(T_0)}}(\mathbf{x}_i)\big) = \operatorname{sign}(y_i)$ holds with high probability. We can also prove for labeled data $(\mathbf{x}_i', y_i')$ that $y_i' \cdot f_{\mathbf{W}^{(T_0)}}(\mathbf{x}_i') \in \big[\big(1-o(1)\big)\cdot(\widehat{\Lambda}_{y_i'}^{(T_0)})^3, \big(m+o(1)\big)\cdot(\widehat{\Lambda}_{y_i'}^{(T_0)})^3\big]$, $\operatorname{sign}\big(f_{\mathbf{W}^{(T_0)}}(\mathbf{x}_i')\big) = \operatorname{sign}(y_i')$ in the same way.

Note that $\widehat{y}_i$ takes y_i with probability p, $-y_i$ with probability p and $n_{\mathrm{l}} = o(n_{\mathrm{u}})$, the first statement in this lemma follows obviously.

To prove the other two statement, we need to give an upper bound for the norm of $\mathbf{w}_j$. According to the update rule of $\mathbf{w}_j^{(t)}$, we have

$$\begin{aligned} \mathbf{w}_j^{(t+1)} = (1-\eta\lambda)\cdot \mathbf{w}_j^{(t)} + \frac{3\eta u_j}{n_{\mathrm{l}}+n_{\mathrm{u}}}\cdot\Bigg(&\sum_{i=1}^{n_{\mathrm{u}}} c_i\widehat{y}_i\big([\langle \mathbf{w}_j, y_i\cdot\mathbf{v}\rangle]_+^2\cdot y_i\cdot\mathbf{v} + [\langle \mathbf{w}_j, \boldsymbol{\xi}_i\rangle]_+^2\cdot\boldsymbol{\xi}_i\big) \\ &+ \sum_{i=1}^{n_{\mathrm{l}}} b_i y_i'\big([\langle \mathbf{w}_j, y_i'\cdot\mathbf{v}\rangle]_+^2\cdot y_i'\cdot\mathbf{v} + [\langle \mathbf{w}_j, \boldsymbol{\xi}_i'\rangle]_+^2\cdot\boldsymbol{\xi}_i'\big)\Bigg), \end{aligned}$$

leading to

$$\begin{aligned} \|\mathbf{w}_j^{(t+1)}\|_2 &\leq (1-\eta\lambda)\cdot\|\mathbf{w}_j^{(t)}\|_2 + \frac{3\eta}{n_{\mathrm{l}}+n_{\mathrm{u}}}\cdot\Bigg(\sum_{i=1}^{n_{\mathrm{u}}}\big([\langle \mathbf{w}_j, y_i\cdot\mathbf{v}\rangle]_+^2\cdot\|\mathbf{v}\|_2 + [\langle \mathbf{w}_j, \boldsymbol{\xi}_i\rangle]_+^2\cdot\|\boldsymbol{\xi}_i\|_2\big) \\ &\qquad + \sum_{i=1}^{n_{\mathrm{l}}}\big([\langle \mathbf{w}_j, y_i'\cdot\mathbf{v}\rangle]_+^2\cdot\|\mathbf{v}\|_2 + [\langle \mathbf{w}_j, \boldsymbol{\xi}_i'\rangle]_+^2\cdot\|\boldsymbol{\xi}_i'\|_2\big)\Bigg) \\ &\leq (1-\eta\lambda)\cdot\|\mathbf{w}_j^{(t)}\|_2 + \frac{3\eta}{n_{\mathrm{l}}+n_{\mathrm{u}}}\cdot\Bigg((n_{\mathrm{l}}+n_{\mathrm{u}})\cdot\|\mathbf{v}\|_2\cdot\big(\max_{r\in\{\pm1\}}\{\widehat{\Lambda}_r^{(t)}, \bar{\Lambda}_r^{(t)}\}\big)^2 \\ &\qquad + \Big(\sum_{i\in[n_{\mathrm{u}}]}\|\boldsymbol{\xi}_i\|_2 + \sum_{i\in[n_{\mathrm{l}}]}\|\boldsymbol{\xi}_i'\|_2\Big)\cdot\big(\Gamma^{(t)}\big)^2\Bigg) \\ &\leq \|\mathbf{w}_j^{(t)}\|_2 + \eta\cdot\Big(\Theta(d^{\frac{1}{2}})\cdot\widetilde{\Theta}(1) + \Theta(d^{\frac{1}{2}+\epsilon})\cdot\widetilde{O}(d^{-\frac{1}{2}+2\epsilon})\Big) \\ &= \|\mathbf{w}_j^{(t)}\|_2 + \eta\cdot\widetilde{\Theta}(d^{\frac{1}{2}}), \end{aligned} \tag{A.43}$$

where the first inequality is by triangle inequality; the second inequality is due to the definition of $\widehat{\Lambda}_r^{(t)}, \bar{\Lambda}_r^{(t)}, \Gamma^{(t)}$, the last inequality is due to lemma 5.3.

According to lemma A.16, we know that $T_r\cdot\eta = \widetilde{\Theta}(d^{-\frac{3}{4}}), r\in\{\pm1\}$ and $T_0\cdot\eta = \max_{r\in\{\pm1\}}\{T_r\cdot\eta\} = \widetilde{\Theta}(d^{-\frac{3}{4}})$. Note that $\mathbf{w}_j^{(0)} \sim \mathcal{N}(\mathbf{0}, \sigma_0^2\mathbf{I}_d), \sigma_0 = \Theta(d^{-\frac{3}{4}})$ and hence $\|\mathbf{w}_j^{(0)}\|_2 = \widetilde{\Theta}(d^{-\frac{1}{4}})$, we know that

$$\|\mathbf{w}_j^{(T_0)}\|_2 \leq \|\mathbf{w}_j^{(0)}\|_2 + \eta\cdot T_0\cdot\widetilde{\Theta}(d^{-\frac{1}{4}}) = \widetilde{\Theta}(d^{-\frac{1}{4}}) + \widetilde{\Theta}(d^{-\frac{1}{4}}) = \widetilde{\Theta}(d^{-\frac{1}{4}}).$$

Therefore, for any $(\mathbf{x}, y)$ sampled from distribution $\mathcal{D}$ where $\mathbf{x} = [y\cdot\mathbf{v}^\top, \boldsymbol{\xi}^\top]^\top$ and $\boldsymbol{\xi}\sim\mathcal{N}(0,\sigma_p^2)$, we have

$$\langle \mathbf{w}_j^{(T_0)}, \boldsymbol{\xi}\rangle \sim \mathcal{N}(0, \sigma_p^2\|\mathbf{w}_j^{(T_0)}\|_2^2), |\langle \mathbf{w}_j^{(T_0)}, \boldsymbol{\xi}\rangle| = \Theta(\sigma_p\|\mathbf{w}_j^{(T_0)}\|_2) = \widetilde{O}(d^{-\frac{1}{4}+\epsilon}).$$

And this indicates that $\langle \mathbf{w}_j^{(T_0)}, \boldsymbol{\xi}\rangle$ will still be dominated by $\langle \mathbf{w}_j^{(T_0)}, \mathbf{v}\rangle$, therefore it holds for newly sampled $(\mathbf{x}, y)$ that

$$y\cdot f_{\mathbf{W}^{(T_0)}}(\mathbf{x}) \in \big[\big(1-o(1)\big)\cdot(\widehat{\Lambda}_{y_i}^{(T_0)})^3, \big(m+o(1)\big)\cdot(\widehat{\Lambda}_{y_i}^{(T_0)})^3\big],$$

which means that

$$\mathbb{P}_{(\mathbf{x},y)\sim\mathcal{D}}[y \cdot f_{\mathbf{W}^{(T_0)}}(\mathbf{x}) \leq 0] = o(1).$$

This verifies the third statement that test error is nearly zero.

For the second statement, note that

$$\begin{aligned}
&\mathbb{P}_{(\mathbf{x},y)\sim\mathcal{D},\widehat{y}\sim y\cdot\mathcal{B}(p)}[\widehat{y} \cdot f_{\mathbf{W}^{(T_0)}}(\mathbf{x}) \leq 0] \\
&= \mathbb{P}_{(\mathbf{x},y)\sim\mathcal{D}}[\widehat{y} \cdot f_{\mathbf{W}^{(T_0)}}(\mathbf{x}) \leq 0 | \widehat{y} = y] \cdot \mathbb{P}_{\widehat{y}\sim y\cdot\mathcal{B}(p)}(\widehat{y} = y) \\
&\qquad + \mathbb{P}_{(\mathbf{x},y)\sim\mathcal{D}}[\widehat{y} \cdot f_{\mathbf{W}^{(T_0)}}(\mathbf{x}) \leq 0 | \widehat{y} = -y] \cdot \mathbb{P}_{\widehat{y}\sim y\cdot\mathcal{B}(p)}(\widehat{y} = -y) \\
&= p \cdot \mathbb{P}_{(\mathbf{x},y)\sim\mathcal{D}}[y \cdot f_{\mathbf{W}^{(T_0)}}(\mathbf{x}) \leq 0] + (1-p) \cdot \mathbb{P}_{(\mathbf{x},y)\sim\mathcal{D}}[y \cdot f_{\mathbf{W}^{(T_0)}}(\mathbf{x}) \geq 0] \\
&= p \cdot o(1) + (1-p) \cdot (1 - o(1)) \\
&= 1 - p \pm o(1),
\end{aligned}$$

which verifies the second statement. □

A.8 Downstream task

For downstream tasks, we use early stopped classifiers, which are stopped when on-diagonal feature $\widehat{\Lambda}_r^{(t)}$ are learned while off-diagonal feature $\bar{\Lambda}_r^{(t)}$ and noise $\Gamma^{(t)}$ are not memorized. Assume we have learned K early stopped classifiers $f_{\mathbf{W}_1^{(T_0^1)}}(\mathbf{x}), \cdots, f_{\mathbf{W}_K^{(T_0^K)}}(\mathbf{x})$ by using n_u pseudo-labeled data generated by pseudolabeler $f_1^{\mathrm{w}}, \cdots, f_K^{\mathrm{w}}$ and n_{l} labeled data.

Then, we want to design a classifier on the learned representation $f_{\mathbf{W}_1^{(T_0^1)}}(\mathbf{x}), \cdots, f_{\mathbf{W}_K^{(T_0^K)}}(\mathbf{x})$ to fit y. Here we consider training a downstream linear model

$$g_{\mathbf{a}}(\mathbf{x}) = \sum_{k=1}^{K} a_k f_{\mathbf{W}_k^{(T_0^k)}}(\mathbf{x}),$$

where $a_k \in \mathbb{R}$ denotes the weight as the k-th pretext model. Given labeled training data $S' = \{(\mathbf{x}_i', y_i')\}_{i=1}^{n_{\mathrm{l}}}$, we want to optimize the empirical loss function

$$L_{S'}(\mathbf{a}) = \frac{1}{n_{\mathrm{l}}} \sum_{i=1}^{n_{\mathrm{l}}} \ell\big(y_i' \cdot g_{\mathbf{a}}(\mathbf{x}_i')\big),$$

where $\ell(z) = \log(1 + \exp(-z))$ denotes the cross entropy loss. We initialize $\mathbf{a}$ as zero and optimize empirical loss function by gradient descent, i.e.

$$\mathbf{a}^{(t+1)} = \mathbf{a}^{(t)} - \eta \cdot \nabla_{\mathbf{a}} L_{S'}(\mathbf{a}^{(t)}), \mathbf{a}^{(0)} = \mathbf{0}.$$

In order to estimate the training error and test error for downstream task, we first introduce following lemma about the increasing rate of $\big\|\mathbf{a}^{(t)}\big\|_1$.

Lemma A.19 (Logarithmic increasing rate)**.** For any learning rate $\eta > 0$, $a_k^{(t)}$ will always increase for any $k \in [K]$ and hence $\big\|\mathbf{a}^{(t)}\big\|_1 = \sum_{k=1}^{K} a_k^{(t)}$. And it holds that $\big\|\mathbf{a}^{(t)}\big\|_1 = \Theta(\log(t))$.

In order to give the increasing rate of $\big\|\mathbf{a}^{(t)}\big\|_1$, we introduce and prove the following lemma:

Lemma A.20. Consider following sequence $\{x_n\}_{n=0}^{\infty}$

$$x_{n+1} = x_n + A \cdot e^{-C \cdot x_n}, x_0 = 0,$$

where A, C are positive constants,and it follows that $x_n = \frac{\log n}{C}\{+ \text{ lower order terms}\}$.

Proof of lemma A.20. Note that

$$x_{n+1} = x_n + A \cdot e^{-C \cdot x_n} \iff e^{C \cdot x_n} \cdot (x_{n+1} - x_n) = \frac{1}{A} \implies \sum_{i=0}^{n-1} e^{C \cdot x_n} \cdot (x_{n+1} - x_n) = \frac{n}{A}.$$

Since x_n and hence $e^{C\cdot x_n}$ are monotone increasing, this means

$$\int_{x_0}^{x_n} e^{C\cdot x} dx \geq \frac{n}{A},$$

implying

$$\frac{e^{C\cdot x}}{C}\bigg|_{x_0}^{x_n} = \frac{e^{C\cdot x_n} - e^{C\cdot x_0}}{C} = \frac{e^{C\cdot x_n} - 1}{C} \geq \frac{n}{A} \text{ and } e^{C\cdot x_n} \geq \frac{C}{A}\cdot n + 1$$

This gives

$$x_n \geq \frac{\log\big((C/A)\cdot n + 1\big)}{C} \text{ and } x_{n+1} - x_n = e^{-C\cdot x_n} \leq \frac{1}{(C/A)\cdot n + 1} \tag{A.44}$$

which indicates that x_n diverges but very slowly. Whatever, we know

$$e^{C\cdot x_{n+1}} = e^{C\cdot x_n}\cdot e^{C\cdot e^{-C\cdot x_n}} \leq e^{C\cdot x_n}\cdot e^{C/(C\cdot n+1)} \leq 3e^{C\cdot x_n},$$

which implies

$$\sum_{k=0}^{n-1} e^{C\cdot x_{n+1}}\cdot(x_{n+1} - x_n) \leq 3\sum_{k=0}^{n-1} e^{C\cdot x_n}\cdot(x_{n+1} - x_n) = \frac{3}{A}\cdot n.$$

Since $e^{C\cdot x_n}$ is monotone increasing, this means

$$\frac{e^{C\cdot x_n} - 1}{C} = \int_{x_0}^{x_n} e^{C\cdot x} dx \leq \sum_{k=0}^{n-1} e^{C\cdot x_{n+1}}\cdot(x_{n+1} - x_n) \leq \frac{3}{A}\cdot n$$

leading to

$$x_n \leq \frac{\log\big((3C/A)\cdot n + 1\big)}{C}. \tag{A.45}$$

From A.44 and A.45, we know that

$$\frac{\log\big((C/A)\cdot n + 1\big)}{C} \leq x_n \leq \frac{\log\big((3C/A)\cdot n + 1\big)}{C}$$

and

$$x_n = \frac{\log n}{C}\{+ \text{ lower order terms}\},$$

which completes the proof. □

Now we are ready to prove lemma A.19.

Proof of lemma A.19. Note that we take downstream task linear model $g_{\mathbf{a}}(\mathbf{x})$ as

$$\begin{aligned} g_{\mathbf{a}}(\mathbf{x}) = \sum_{k=1}^{d} a_k \Bigg\{ &\sum_{j=1}^{m} \Big[\sigma\big(\langle \mathbf{w}_{k,j}^{(T_0^k)}, y\cdot\mathbf{v}\rangle\big) + \sigma\big(\langle \mathbf{w}_{k,j}^{(T_0^k)}, \boldsymbol{\xi}\rangle\big)\Big] \\ &- \sum_{j=m+1}^{2m} \Big[\sigma\big(\langle \mathbf{w}_{k,j}^{(T_0^k)}, y\cdot\mathbf{v}\rangle\big) + \sigma\big(\langle \mathbf{w}_{k,j}^{(T_0^k)}, \boldsymbol{\xi}\rangle\big)\Big] \Bigg\} \\ = \sum_{k=1}^{d} a_k f_{\mathbf{W}_k^{(T_0^k)}}(\mathbf{x}). \end{aligned}$$

Then, we have following update rule for model parameter $\mathbf{a}$:

$$a_k^{(t+1)} = a_k^{(t)} - \eta\cdot\frac{1}{n_1}\sum_{i=1}^{n_1} \ell'\big(y_i'\cdot g_{\mathbf{a}^{(t)}}(\mathbf{x}_i')\big)\cdot y_i' f_{\mathbf{W}_k^{(T_0^k)}}(\mathbf{x}_i'),$$

where we initialize $a_k^{(0)}$ as zero for all $k \in [K]$.

Next, we prove following statement by using induction method: when $t \geq 1$,

- $a_k^{(t)}, \forall k \in [K]$ is non-negative and increasing.
- $\left\|\mathbf{a}^{(t)}\right\|_1 = \sum_{i=1}^K a_k^{(t)}$.
- $a_k^{(t+1)} = a_k^{(t)} + \eta \cdot \widetilde{\Theta}(1) \cdot \Big(\exp\big(- \|\mathbf{a}^{(1)}\|_1 \cdot \widetilde{\Theta}(1)\big)\Big), \forall k \in [K]$.

Note that $a_k^{(0)} = 0$ for all $k \in [d]$ and therefore $g_{\mathbf{a}^{(0)}}(\mathbf{x}_i') = 0$, $\ell'\big(y_i' \cdot g_{\mathbf{a}^{(0)}}(\mathbf{x}_i')\big) = \ell'(0) = -1/2$,

$$
\begin{aligned}
a_k^{(1)} &= a_k^{(0)} - \eta \cdot \frac{1}{n_1} \sum_{i=1}^{n_1} \ell'\big(y_i' \cdot g_{\mathbf{a}^{(0)}}(\mathbf{x}_i')\big) \cdot y_i' f_{\mathbf{W}_k^{(T_0^k)}}(\mathbf{x}_i') \\
&= a_k^{(0)} + \eta \cdot \frac{1}{2n_1} \sum_{i=1}^{n_1} y_i' f_{\mathbf{W}_k^{(T_0^k)}}(\mathbf{x}_i') = \eta \cdot \frac{1}{2n_1} \sum_{i=1}^{n_1} y_i' f_{\mathbf{W}_k^{(T_0^k)}}(\mathbf{x}_i') \text{ for all } k \in [K].
\end{aligned}
$$

Note that the accuracy of the k-th pseudolabeler $p_k > 1/2$, accoring to the proof of lemma A.18, we have

$$
\begin{aligned}
f_{\mathbf{W}_k^{(T_0^k)}}(\mathbf{x}_i') &= \sum_{j=1}^{m} \Big[\sigma\big(\langle \mathbf{w}_{k,j}^{(T_0^k)}, y_i' \cdot \mathbf{v}\rangle\big) + \sigma\big(\langle \mathbf{w}_{k,j}^{(T_0^k)}, \boldsymbol{\xi}_i'\rangle\big)\Big] \\
&\qquad - \sum_{j=m+1}^{2m} \Big[\sigma\big(\langle \mathbf{w}_{k,j}^{(T_0^k)}, y_i' \cdot \mathbf{v}\rangle\big) + \sigma\big(\langle \mathbf{w}_{k,j}^{(T_0^k)}, \boldsymbol{\xi}_i'\rangle\big)\Big] \\
&= y_i' \cdot \widetilde{\Theta}\big((\widehat{\Lambda}_{y_i'}^{(T_0^k)})^3\big),
\end{aligned}
$$

for all $k \in [K]$. Therefore

$$
a_k^{(1)} = \eta \cdot \frac{1}{2n_1} \sum_{i=1}^{n_1} y_i' f_{\mathbf{W}_k^{(T_0^k)}}(\mathbf{x}_i') \geq \frac{\eta}{2} \cdot \widetilde{\Theta}\big((\widehat{\Lambda}_{y_i'}^{(T_0^k)})^3\big) > 0, \forall k \in [K].
$$

It follows that

$$
\left\|\mathbf{a}^{(t)}\right\|_1 = \sum_{i=1}^{K} |a_k^{(t)}| = \sum_{i=1}^{K} a_k^{(t)}.
$$

Note that

$$
\begin{aligned}
y_i' \cdot g_{\mathbf{a}^{(1)}}(\mathbf{x}_i') &= y_i' \cdot \sum_{k=1}^{d} a_k^{(1)} f_{\mathbf{W}_k^{(T_0^k)}}(\mathbf{x}_i') \\
&= \sum_{k=1}^{d} a_k^{(1)} \cdot \Big(y_i' \cdot f_{\mathbf{W}_k^{(T_0^k)}}(\mathbf{x}_i')\Big) \\
&= \sum_{k=1}^{d} a_k^{(1)} \cdot \widetilde{\Theta}\big((\widehat{\Lambda}_{y_i'}^{(T_0^k)})^3\big) \\
&= \sum_{k=1}^{d} a_k^{(1)} \cdot \widetilde{\Theta}(1) \\
&= \|\mathbf{a}^{(1)}\|_1 \cdot \widetilde{\Theta}(1).
\end{aligned} \tag{A.46}
$$

This leads to

$$
\begin{aligned}
\ell'\big(y_i' \cdot g_{\mathbf{a}^{(1)}}(\mathbf{x}_i')\big) &= -\frac{\exp\big(-y_i' \cdot g_{\mathbf{a}^{(1)}}(\mathbf{x}_i')\big)}{1 + \exp\big(-y_i' \cdot g_{\mathbf{a}^{(1)}}(\mathbf{x}_i')\big)} \\
&= -c \cdot \Big(\exp\big(- y_i' \cdot g_{\mathbf{a}^{(1)}}(\mathbf{x}_i')\big)\Big) \\
&= -c \cdot \Big(\exp\big(- \|\mathbf{a}^{(1)}\|_1 \cdot \widetilde{\Theta}(1)\big)\Big),
\end{aligned}
$$

where the second equality is due to $y_i' \cdot g_{\mathbf{a}^{(1)}}(\mathbf{x}_i') > 0$, $\exp\left(-y_i' \cdot g_{\mathbf{a}^{(1)}}(\mathbf{x}_i')\right) < 1$ and $c \in (1/2, 1)$; the last equality is due to (A.46). It follows that

$$\begin{aligned} a_k^{(2)} &= a_k^{(1)} - \eta \cdot \frac{1}{n_l} \sum_{i=1}^{n_l} \ell'\big(y_i' \cdot g_{\mathbf{a}^{(1)}}(\mathbf{x}_i')\big) \cdot y_i' f_{\mathbf{W}_k^{(T_0)}}(\mathbf{x}_i') \\ &= a_k^{(1)} + \eta \cdot c \cdot \widetilde{\Theta}(1) \cdot \Big(\exp\big(-\|\mathbf{a}^{(1)}\|_1 \cdot \widetilde{\Theta}(1)\big)\Big), \forall k \in [K] \end{aligned}$$

where $c \in (1/2, 1)$. By then, we have already proved the induction hypothesis of $t = 1$.

Next, assume the induction hypotheses hold for t. For $t+1$, we have

$$a_k^{(t+1)} = a_k^{(t)} - \eta \cdot \frac{1}{n_l} \sum_{i=1}^{n_l} \underbrace{\ell'\big(y_i' \cdot g_{\mathbf{a}^{(t)}}(\mathbf{x}_i')\big)}_{<0} \cdot \underbrace{y_i' f_{\mathbf{W}_k^{(T_0^k)}}(\mathbf{x}_i')}_{>0} > a_k^{(t)} > 0.$$

And it follows that

$$\|\mathbf{a}^{(t+1)}\|_1 = \sum_{i=1}^{K} a_k^{(t+1)} \text{ and } y_i' \cdot g_{\mathbf{a}^{(t+1)}}(\mathbf{x}_i') = \|\mathbf{a}^{(t+1)}\|_1 \cdot \widetilde{\Theta}(1), \tag{A.47}$$

leading to

$$\ell'\big(y_i' \cdot g_{\mathbf{a}^{(t+1)}}(\mathbf{x}_i')\big) = -c \cdot \Big(\exp\big(-\|\mathbf{a}^{(t+1)}\|_1 \cdot \widetilde{\Theta}(1)\big)\Big), c \in (1/2, 1),$$

and

$$a_k^{(t+2)} = a_k^{(t+1)} + \eta \cdot \widetilde{\Theta}(1) \cdot \Big(\exp\big(-\|\mathbf{a}^{(t+1)}\|_1 \cdot \widetilde{\Theta}(1)\big)\Big), \forall k \in [K].$$

This indicates that if induction hypotheses hold for t, then they holds for $t+1$.

Adding up $k \in [K]$, we can obtain

$$\|\mathbf{a}^{(t+1)}\|_1 = \left\|\mathbf{a}^{(t)}\right\|_1 + \eta \cdot \widetilde{\Theta}(1) \cdot \exp\Big(-\widetilde{\Theta}(1) \cdot \left\|\mathbf{a}^{(t)}\right\|_1\Big)$$

According to lemma (A.20), we know that $\left\|\mathbf{a}^{(t)}\right\|_1 = \log t / \widetilde{\Theta}(1)\{\pm \text{ lower order terms w.r.t. t}\}$. □

The following lemma gives the convergence guarantee of downstream task:

Lemma A.21. (Convergence Guarantee) For any learning rate $\eta > 0$,

$$\|\nabla_{\mathbf{a}} L_S(\mathbf{a}^{(t)})\|_1 = \exp\big(-\left\|\mathbf{a}^{(t)}\right\|_1 \cdot \widetilde{\Theta}(1)\big) \text{ and } \nabla_{\mathbf{a}}^2 L_S(\mathbf{a}) \succeq 0 \text{ for any } \mathbf{a} \in \mathbb{R}^d.$$

Proof of lemma A.21. Recall the partial derivative that

$$\partial_{a_k} L_{S'}(\mathbf{a}) = \frac{1}{n_l} \sum_{i=1}^{n_l} \ell'\big(y_i' \cdot g_{\mathbf{a}^{(t)}}(\mathbf{x}_i')\big) \cdot y_i' f_{\mathbf{W}_k^{(T_0)}}(\mathbf{x}_i'),$$

leading to

$$\begin{aligned} \|\nabla_{\mathbf{a}} L_{S'}(\mathbf{a}^{(t)})\|_1 &= \sum_{k=1}^{K} \left| \frac{1}{n_l} \sum_{i=1}^{n_l} \ell'\big(y_i' \cdot g_{\mathbf{a}^{(t)}}(\mathbf{x}_i')\big) \cdot y_i' f_{\mathbf{W}_k^{(T_0^k)}}(\mathbf{x}_i') \right| \\ &= \sum_{k=1}^{K} \frac{1}{n_l} \sum_{i=1}^{n_l} \left|\ell'\big(y_i' \cdot g_{\mathbf{a}^{(t)}}(\mathbf{x}_i')\big)\right| \cdot \left|y_i' f_{\mathbf{W}_k^*}(\mathbf{x}_i')\right| \\ &= \frac{1}{n_l} \sum_{i=1}^{n_l} \left|\ell'\big(y_i' \cdot g_{\mathbf{a}^{(t)}}(\mathbf{x}_i')\big)\right| \cdot y_i' f_{\mathbf{W}_k^*}(\mathbf{x}_i') \\ &= \frac{1}{n_l} \sum_{i=1}^{n_l} \left|\ell'\big(y_i' \cdot g_{\mathbf{a}^{(t)}}(\mathbf{x}_i')\big)\right| \cdot \widetilde{\Theta}(1). \end{aligned}$$

Note that

$$\begin{aligned}\left|\ell'\big(y_i' \cdot g_{\mathbf{a}^{(t)}}(\mathbf{x}_i')\big)\right| &= \frac{\exp\left(-y_i' \cdot g_{\mathbf{a}^{(t)}}(\mathbf{x}_i')\right)}{1+\exp\left(-y_i' \cdot g_{\mathbf{a}^{(t)}}(\mathbf{x}_i')\right)} \\ &= \Theta\Big(\exp\big(-y_i' \cdot g_{\mathbf{a}^{(t)}}(\mathbf{x}_i')\big)\Big) \\ &= \Theta\Big(\exp\big(-\left\|\mathbf{a}^{(t)}\right\|_1 \cdot \widetilde{\Theta}(1)\big)\Big),\end{aligned}$$

it follows that

$$\|\nabla_{\mathbf{a}} L_{S'}(\mathbf{a}^{(t)})\|_1 = \exp\big(-\left\|\mathbf{a}^{(t)}\right\|_1 \cdot \widetilde{\Theta}(1)\big) = \exp\big(-\Theta(\log t)\cdot\widetilde{\Theta}(1)\big) = \Theta\big(1/\mathrm{poly}(t)\big),$$

which shows that within polynomial steps, gradient descent is guaranteed to find a point with small gradient.

Note that

$$\partial_{a_k}\partial_{a_j} L_{S'}(\mathbf{a}) = \frac{1}{n_1}\sum_{i=1}^{n_1} \ell''\big(y_i' \cdot g_{\mathbf{a}^{(t)}}(\mathbf{x}_i')\big)\cdot\Big(f_{\mathbf{W}_k^{(T_0^k)}}(\mathbf{x}_i)\cdot f_{\mathbf{W}_j^{(T_0^j)}}(\mathbf{x}_i)\Big) \text{ for all } k,j\in[K],$$

Denote $\Big[f_{\mathbf{W}_1^{(T_0^1)}}(\mathbf{x}_i'),\cdots,f_{\mathbf{W}_K^{(T_0^K)}}(\mathbf{x}_i')\Big]^\top$ as $\mathbf{f}_{\mathbf{W}^*}(\mathbf{x}_i')$, then

$$\nabla^2_{\mathbf{a}} L_S(\mathbf{a}) = \frac{1}{n_1}\sum_{i=1}^{n_1}\ell''\big(y_i' \cdot g_{\mathbf{a}^{(t)}}(\mathbf{x}_i')\big)\cdot\Big(\mathbf{f}_{\mathbf{W}^*}(\mathbf{x}_i')\cdot\mathbf{f}_{\mathbf{W}^*}(\mathbf{x}_i')^\top\Big).$$

Note that $\mathbf{f}_{\mathbf{W}^*}(\mathbf{x}_i')\cdot\mathbf{f}_{\mathbf{W}^*}(\mathbf{x}_i')^\top$ is a non-negative definite matrix, $\ell''(z) = \exp(-z)/\big(1+\exp(-z)\big)^2 > 0$ and the fact that sum of non-negative definite matrices is still a non-negative definite matrix, it follows that $\nabla^2_{\mathbf{a}} L_S(\mathbf{a}) \succeq 0$. □

Theorem A.22 (Restatement of theorem 4.3)**.** Under self-supervised learning setting, for downstream task, suppose K early stopped classifiers $\{f_{\mathbf{W}_k^*}\}_{k=1}^K$ are obtained after the training of K pretext tasks finished, and after $T_{\mathrm{dt}} = \widetilde{\Theta}(1)$ iterations with learning rate $\eta > 0$, then we can find a linear model $\mathbf{a}^{(T_{\mathrm{dt}})}$, which satisfies: Both test error and loss are nearly 0, i.e. $\mathbb{P}_{(\mathbf{x},y)\sim\mathcal{D}}[y\cdot g_{\mathbf{a}^{(T_{\mathrm{dt}})}}(\mathbf{x}) \le 0] = o(1), L_{\mathcal{D}}(\ell(y\cdot g_{\mathbf{a}^{(T_{\mathrm{dt}})}}(\mathbf{x}))) = o(1)$.

Proof of theorem A.22. For test error, we have

$$\begin{aligned}\mathbb{P}_{(\mathbf{x},y)\sim\mathcal{D}}[y\cdot g_{a^{(T_{\mathrm{dt}})}}(\mathbf{x})\le 0] &= \mathbb{P}_{(\mathbf{x},y)\sim\mathcal{D}}\left[\sum_{k=1}^K a_k^{(T_{\mathrm{dt}})}\cdot\big(y\cdot f_{\mathbf{W}_k^*}(\mathbf{x})\big)\le 0\right] \\ &= \mathbb{P}_{(\mathbf{x},y)\sim\mathcal{D}}\left[\sum_{k=1}^K a_k^{(T_{\mathrm{dt}})}\cdot\widetilde{\Theta}(1)\le 0\right] = o(1)\end{aligned}$$

where the last equality is due to $a_k^{(T_{\mathrm{dt}})} > 0$ according to lemma A.19.

For test loss, we have

$$\begin{aligned}L_{\mathcal{D}}(\ell(y\cdot g_{\mathbf{a}^{(T_{\mathrm{dt}})}}(\mathbf{x}))) &= \mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}[\ell(y\cdot g_{\mathbf{a}^{(T_{\mathrm{dt}})}}(\mathbf{x}))] \\ &= \mathbb{E}_{(\mathbf{x},y)\sim\mathcal{D}}\big[\ell\big(\|\mathbf{a}\|_1^{(T_{\mathrm{dt}})}\cdot\widetilde{\Theta}(1)\big)\big] + o(1) \\ &= \ell\big(\|\mathbf{a}\|_1^{(T_{\mathrm{dt}})}\cdot\widetilde{\Theta}(1)\big) + o(1) \\ &= \ell\big(\log(t)\cdot\widetilde{\Theta}(1)\big) + o(1) = \log\big(1+t^{-\widetilde{\Theta}(1)}\big) + o(1) \\ &= t^{-\widetilde{\Theta}(1)}\{+\text{ lower order terms w.r.t. } t\} + o(1),\end{aligned}$$

where the second equality is due to (A.46) and (A.47); the fourth equality is by lemma A.19; the last equality is due to Taylor expansion. This shows that $L_{\mathcal{D}}(\ell(y\cdot g_{\mathbf{a}^{(t)}}(\mathbf{x})))$ approaches zero at a polynomial rate. Therefore, as long as $t = \Omega(\log^\alpha(d)), \alpha > 0$, we have $t^{-\widetilde{\Theta}(1)} = o(1)$ and it follows that $L_{\mathcal{D}}(\ell(y\cdot g_{\mathbf{a}^{(t)}}(\mathbf{x}))) = o(1)$. □

B PROOF OF SUPERVISED LEARNING SETTING

Here we prove theorem 4.4. First, we give following lemma to facilitate the proof.

Lemma B.1 (Gradient Calculation)**.** The gradient of loss function $L_S(\mathbf{W})$ with respect to weight parameter $\mathbf{w}_j$ is

$$\nabla_{\mathbf{w}_j} L_{S'}(\mathbf{W}) = -\frac{3u_j}{n_l} \cdot \sum_{i=1}^{n_l} b_i y_i' \big([\langle \mathbf{w}_j, y_i' \cdot \mathbf{v} \rangle]_+^2 \cdot y_i' \cdot \mathbf{v} + [\langle \mathbf{w}_j, \boldsymbol{\xi}_i' \rangle]_+^2 \cdot \boldsymbol{\xi}_i' \big) + \lambda \cdot \mathbf{w}_j,$$

where $u_j := \big(\mathbb{1}_{[1 \le j \le m]} - \mathbb{1}_{[m+1 \le j \le 2m]} \big)$ and $-\ell'\big(y_i' \cdot f_{\mathbf{W}}(\mathbf{x}_i')\big) = \exp[-y_i' \cdot f_{\mathbf{W}}(\mathbf{x}_i')]/(1 + \exp[-y_i' \cdot f_{\mathbf{W}}(\mathbf{x}_i')])$ is denoted as b_i.

Proof of Lemma B.1. When $1 \le j \le m$,

$$\begin{aligned}
\nabla_{\mathbf{w}_j} \ell\big(y_i' \cdot f_{\mathbf{W}}(\mathbf{x}_i')\big) &= \ell'\big(y_i' \cdot f_{\mathbf{W}}(\mathbf{x}_i')\big) \cdot y_i' \cdot \nabla_{\mathbf{w}_j} f_{\mathbf{W}}(\mathbf{x}_i') \\
&= -b_i \cdot y_i' \cdot \nabla_{\mathbf{w}_j} f_{\mathbf{W}}(\mathbf{x}_i') \\
&= -b_i y_i' \cdot \big(\sigma'\big(\langle \mathbf{w}_j, y_i' \cdot \mathbf{v} \rangle\big) \cdot y_i' \cdot \mathbf{v} + \sigma'\big(\langle \mathbf{w}_j, \boldsymbol{\xi}_i' \rangle\big) \cdot \boldsymbol{\xi}_i' \big) \\
&= -3 b_i y_i' \big([\langle \mathbf{w}_j, y_i' \cdot \mathbf{v} \rangle]_+^2 \cdot y_i' \cdot \mathbf{v} + [\langle \mathbf{w}_j, \boldsymbol{\xi}_i' \rangle]_+^2 \cdot \boldsymbol{\xi}_i' \big)
\end{aligned}$$

and when $m+1 \le j \le 2m$,

$$\nabla_{\mathbf{w}_j} \ell\big(y_i' \cdot f_{\mathbf{W}}(\mathbf{x}_i')\big) = 3 b_i y_i' \big([\langle \mathbf{w}_j, y_i' \cdot \mathbf{v} \rangle]_+^2 \cdot y_i' \cdot \mathbf{v} + [\langle \mathbf{w}_j, \boldsymbol{\xi}_i' \rangle]_+^2 \cdot \boldsymbol{\xi}_i' \big)$$

Combining above two cases, we have

$$\begin{aligned}
\nabla_{\mathbf{w}_j} \ell\big(y_i' \cdot f_{\mathbf{W}}(\mathbf{x}_i')\big) &= -3\big(\mathbb{1}_{[1 \le j \le m]} - \mathbb{1}_{[m+1 \le j \le 2m]} \big) b_i y_i' \big([\langle \mathbf{w}_j, y_i \cdot \mathbf{v} \rangle]_+^2 \cdot y_i' \cdot \mathbf{v} + [\langle \mathbf{w}_j, \boldsymbol{\xi}_i' \rangle]_+^2 \cdot \boldsymbol{\xi}_i' \big) \\
&= -3 u_j b_i y_i' \big([\langle \mathbf{w}_j, y_i' \cdot \mathbf{v} \rangle]_+^2 \cdot y_i' \cdot \mathbf{v} + [\langle \mathbf{w}_j, \boldsymbol{\xi}_i' \rangle]_+^2 \cdot \boldsymbol{\xi}_i' \big)
\end{aligned}$$

and therefore

$$\begin{aligned}
\nabla_{\mathbf{w}_j} L_{S'}(\mathbf{W}) &= \frac{1}{n_l} \sum_{i=1}^{n_l} \nabla_{\mathbf{w}_j} L_i(\mathbf{W}) + \frac{\lambda}{2} \nabla_{\mathbf{w}_j} \|\mathbf{W}\|_F^2 \\
&= \frac{1}{n_l} \sum_{i=1}^{n_l} \nabla_{\mathbf{w}_j} \ell\big(y_i' \cdot f_{\mathbf{W}}(\mathbf{x}_i')\big) + \frac{\lambda}{2} \nabla_{\mathbf{w}_j} \|\mathbf{W}\|_F^2 \\
&= -\frac{3u_j}{n_l} \sum_{i=1}^{n_l} b_i y_i' \big([\langle \mathbf{w}_j, y_i' \cdot \mathbf{v} \rangle]_+^2 \cdot y_i' \cdot \mathbf{v} + [\langle \mathbf{w}_j, \boldsymbol{\xi}_i' \rangle]_+^2 \cdot \boldsymbol{\xi}_i' \big) + \lambda \cdot \mathbf{w}_j.
\end{aligned}$$

□

Proof of theorem 4.4. Recall the definition of $f_{\mathbf{W}}$ in (3.1) that

$$f_{\mathbf{W}}(\mathbf{x}) = \sum_{j=1}^{m} \Big[\sigma\big(\langle \mathbf{w}_j, y \cdot \mathbf{v} \rangle\big) + \sigma\big(\langle \mathbf{w}_j, \boldsymbol{\xi} \rangle\big) \Big] - \sum_{j=m+1}^{2m} \Big[\sigma\big(\langle \mathbf{w}_j, y \cdot \mathbf{v} \rangle\big) + \sigma\big(\langle \mathbf{w}_j, \boldsymbol{\xi} \rangle\big) \Big].$$

Define $\widetilde{\mathbf{w}}_j := m^{1/3} \cdot \mathbf{w}_j$, we have

$$\begin{aligned}
f_{\mathbf{W}}(\mathbf{x}) &= \sum_{j=1}^{m} \Big[\sigma\big(\langle m^{-1/3} \cdot \widetilde{\mathbf{w}}_j, y \cdot \mathbf{v} \rangle\big) + \sigma\big(\langle m^{-1/3} \cdot \widetilde{\mathbf{w}}_j, \boldsymbol{\xi} \rangle\big) \Big] \\
&\quad - \sum_{j=m+1}^{2m} \Big[\sigma\big(\langle m^{-1/3} \cdot \widetilde{\mathbf{w}}_j, y \cdot \mathbf{v} \rangle\big) + \sigma\big(\langle m^{-1/3} \cdot \widetilde{\mathbf{w}}_j, \boldsymbol{\xi} \rangle\big) \Big] \\
&= \frac{1}{m} \sum_{j=1}^{m} \Big[\sigma\big(\langle \widetilde{\mathbf{w}}_j, y \cdot \mathbf{v} \rangle\big) + \sigma\big(\langle \widetilde{\mathbf{w}}_j, \boldsymbol{\xi} \rangle\big) \Big] - \frac{1}{m} \sum_{j=m+1}^{2m} \Big[\sigma\big(\langle \widetilde{\mathbf{w}}_j, y \cdot \mathbf{v} \rangle\big) + \sigma\big(\langle \widetilde{\mathbf{w}}_j, \boldsymbol{\xi} \rangle\big) \Big] \\
&:= f_{\widetilde{\mathbf{W}}}(\mathbf{x}).
\end{aligned}$$

Since the standard deviation of Gaussian initialization of $\mathbf{w}_j$ is σ_0 and note that $\widetilde{\mathbf{w}}_j := m^{1/3} \cdot \mathbf{w}_j$, the standard deviation of Gaussian initialization of $\widetilde{\mathbf{w}}_j$ is $m^{1/3}\sigma_0 := \widetilde{\sigma}_0$.

On the other hand, note that the update rule of $\mathbf{w}_j^{(t)}$ is $\mathbf{w}_j^{(t+1)} = \mathbf{w}_j^{(t)} - \eta \cdot \nabla_{\mathbf{w}_j} L_{S'}(\mathbf{W}^{(t)})$, and in lemma B.1, we have

$$\nabla_{\mathbf{w}_j} L_{S'}(\mathbf{W}) = -\frac{3u_j}{n_l} \cdot \sum_{i=1}^{n_l} b_i y_i' \big([\langle \mathbf{w}_j, y_i' \cdot \mathbf{v} \rangle]_+^2 \cdot y_i' \cdot \mathbf{v} + [\langle \mathbf{w}_j, \boldsymbol{\xi}_i' \rangle]_+^2 \cdot \boldsymbol{\xi}_i' \big) + \lambda \cdot \mathbf{w}_j.$$

It follows that

$$\mathbf{w}_j^{(t+1)} = \mathbf{w}_j^{(t)} + \frac{3\eta u_j}{n_l} \cdot \sum_{i=1}^{n_l} b_i^{(t)} y_i' \big([\langle \mathbf{w}_j^{(t)}, y_i' \cdot \mathbf{v} \rangle]_+^2 \cdot y_i' \cdot \mathbf{v} + [\langle \mathbf{w}_j^{(t)}, \boldsymbol{\xi}_i' \rangle]_+^2 \cdot \boldsymbol{\xi}_i' \big) + \lambda \cdot \mathbf{w}_j^{(t)}. \quad \text{(B.1)}$$

By plugging $\mathbf{w}_j = m^{-1/3} \cdot \widetilde{\mathbf{w}}_j$ into B.1, we have

$$\widetilde{\mathbf{w}}_j^{(t+1)} = \widetilde{\mathbf{w}}_j^{(t)} + \frac{3\eta m^{-\frac{1}{3}} u_j}{n_l} \cdot \sum_{i=1}^{n_l} b_i^{(t)} y_i' \big([\langle \widetilde{\mathbf{w}}_j^{(t)}, y_i' \cdot \mathbf{v} \rangle]_+^2 \cdot y_i' \cdot \mathbf{v} + [\langle \widetilde{\mathbf{w}}_j^{(t)}, \boldsymbol{\xi}_i' \rangle]_+^2 \cdot \boldsymbol{\xi}_i' \big) + \lambda \cdot \widetilde{\mathbf{w}}_j^{(t)}.$$

Assume $\widetilde{\eta} = m^{-\frac{1}{3}}\eta$, we have $\widetilde{\mathbf{w}}_j^{(t+1)} = \widetilde{\mathbf{w}}_j^{(t)} - \widetilde{\eta} \cdot \nabla_{\widetilde{\mathbf{w}}_j} L_{S'}(\widetilde{\mathbf{W}}^{(t)})$. Therefore, our data model and training algorithm is equivalent to the model and algorithm below:

$$f_{\widetilde{\mathbf{W}}^{+1}}(\mathbf{x}) = \frac{1}{m} \sum_{j=1}^{m} \Big[\sigma\big(\langle \widetilde{\mathbf{w}}_j, y \cdot \mathbf{v} \rangle\big) + \sigma\big(\langle \widetilde{\mathbf{w}}_j, \boldsymbol{\xi} \rangle\big) \Big],$$

$$f_{\widetilde{\mathbf{W}}^{-1}}(\mathbf{x}) = \frac{1}{m} \sum_{j=m+1}^{2m} \Big[\sigma\big(\langle \widetilde{\mathbf{w}}_j, y \cdot \mathbf{v} \rangle\big) + \sigma\big(\langle \widetilde{\mathbf{w}}_j, \boldsymbol{\xi} \rangle\big) \Big],$$

$$f_{\widetilde{\mathbf{W}}}(\mathbf{x}) = f_{\widetilde{\mathbf{W}}^{+1}}(\mathbf{x}) - f_{\widetilde{\mathbf{W}}^{-1}}(\mathbf{x}),$$

and we use gradient decent with learning rate $\widetilde{\eta}$ and cross-entropy loss to optimize such a data model, i.e.

$$\widetilde{\mathbf{w}}_0^{(t)} \sim \mathcal{N}(\mathbf{0}, \widetilde{\sigma}_0^2 \boldsymbol{I}_d), \widetilde{\mathbf{w}}_j^{(t+1)} = \widetilde{\mathbf{w}}_j^{(t)} - \widetilde{\eta} \cdot \nabla_{\widetilde{\mathbf{w}}_j} L_{S'}(\widetilde{\mathbf{W}}^{(t)}),\ L_{S'}(\widetilde{\mathbf{W}}^{(t)}) = \sum_{i=1}^{n_l} \ell(y_i' \cdot f_{\widetilde{\mathbf{W}}}(\mathbf{x}_i')),$$

where $\ell(z) = \log(1 + \exp(-z))$, $\widetilde{\sigma}_0 = m^{1/3}\sigma_0$. Note that the new model meets the one used in Cao et al. (2022). To leverage their result, we introduce condition 4.3 from Cao et al. (2022) and verify that the new model meets the new condition. Notice that our model has activation function $\sigma(z) = [z]_+^3$ which is a special case of $\sigma(z) = [z]_+^q$ considered in Cao et al. (2022) for $q = 3$. We have following simplified condition and Theorems,

Condition B.2 (Condition 4.2 in Cao et al. (2022))**.** Dimension $d = \widetilde{\Omega}(m^4 n^4)$. Training sample size n and neural network width m satisfy $n, m = \Omega(\mathrm{polylog}(d))$. Learning rate η satisfies $\eta \leq \widetilde{O}(\min\{\|\mathbf{v}\|_2^{-2}, \sigma_p^{-2} d^{-1}\})$. The standard deviation of Gaussian initialization σ_0 is approximately chosen such that $\widetilde{O}(n d^{-\frac{1}{2}}) \cdot \min\{(\sigma_p\sqrt{d})^{-1}, \|\mathbf{v}\|_2^{-1}\} \leq \sigma_0 \leq \widetilde{O}(m^{-2} n^{-1}) \cdot \min\{(\sigma_p\sqrt{d})^{-1}, \|\mathbf{v}\|_2^{-1}\}$.

Theorem B.3 (Theorem 4.4 in Cao et al. (2022))**.** For any $\epsilon > 0$, let $T = \widetilde{\Theta}(\eta^{-1} m \cdot n(\sigma_p\sqrt{d})^{-3} \cdot \sigma_0^{-1} + \eta^{-1}\epsilon^{-1} n m^3 d^{-1} \sigma_p^{-2})$. Under condition B.2, if $n^{-1} \cdot \mathrm{SNR}^{-3} = \widetilde{\Omega}(1), \mathrm{SNR} = \|\mathbf{v}\|_2/\sigma_p\sqrt{d}$, then with probability at least $1 - d^{-1}$, there exists $0 \leq t \leq T$ such that:

1. The training loss converges to ϵ, i.e., $L_S(\mathbf{W}^{(t)}) \leq \epsilon$.

2. The trained CNN has a constant order test loss: $L_{\mathcal{D}}(\mathbf{W}^{(t)}) = \Theta(1)$.

Note that in our setting, $m = \Theta(\mathrm{polylog}(d))$, $n_l = \widetilde{\Theta}(1)$, $\|\mathbf{v}\|_2 = \Theta(d^{\frac{1}{2}})$, $\widetilde{\sigma}_0 = m^{1/3}\sigma_0$, $\sigma_0 = \Theta(d^{-\frac{3}{4}})$ $\sigma_p = \Theta(d^{0.01})$, $\widetilde{\eta} = m^{-\frac{1}{3}}\eta$ and $\eta = O(d^{-1.1})$, it's not difficult to verify that condition B.2

holds. Besides, $\mathrm{SNR} = d^{-0.01}, n^{-1} \cdot \mathrm{SNR}^{-q} = \widetilde{\Theta}(d^{0.03}) = \widetilde{\Omega}(1)$. Therefore, the conclusion of Theorem B.3 holds for

$$\begin{aligned} T &= \widetilde{\Theta}(\widetilde{\eta}^{-1} m \cdot n(\sigma_p \sqrt{d})^{-q} \cdot \sigma_0^{-(q-2)} + \widetilde{\eta}^{-1} \epsilon^{-1} n m^3 d^{-1} \sigma_p^{-2}) \\ &= \widetilde{\Theta}(\widetilde{\eta}^{-1} \cdot (d^{0.51})^{-3} \cdot (d^{-0.75})^{-1} + \widetilde{\eta}^{-1} \epsilon^{-1} d^{-1} d^{-0.02}) \\ &= \widetilde{\Theta}(\widetilde{\eta}^{-1} \cdot d^{-0.78}) \\ &= \widetilde{\Theta}(\eta^{-1} \cdot d^{-0.78}). \end{aligned}$$

□

C AUXILIARY LEMMAS

For the estimation of $\|\boldsymbol{\xi}_i\|_2^2$ and $\langle \boldsymbol{\xi}_i, \boldsymbol{\xi}_l \rangle$, we introduce following lemma.

Lemma C.1 (Lemma B.2 in Cao et al. (2022))**.** Suppose that $\delta > 0$ and $d = \Omega(\log(4n/\delta))$. Then with probability at least $1 - \delta$,

$$\begin{aligned} &\sigma_p^2 d/2 \le \|\boldsymbol{\xi}_i\|_2^2 \le 3\sigma_p^2 d/2, \\ &|\langle \boldsymbol{\xi}_i, \boldsymbol{\xi}_l \rangle| \le 2\sigma_p^2 \cdot \sqrt{d \log(4n^2/\delta)}, \end{aligned}$$

for all $i, l \in [n]$, $i \neq l$.

Besides, we introduce following lemma about tensor power method.

Lemma C.2. Consider an increasing sequence $x_t \ge 0$ defined as $x_{t+1} = x_t + \eta \cdot C_t x_t^2$, then we have for $A > x_0$, every $\delta > 0$, and every $\eta > 0$:

$$\sum_{t \ge 0, x_t \le A} \eta \cdot C_t \le \frac{1+\delta}{x_0} + \eta(1+\delta)^2 \Big(1 + \frac{\log(A/x_0)}{\log(1+\delta)}\Big),$$

$$\sum_{t \ge 0, x_t \le A} \eta \cdot C_t \ge \frac{1 - \frac{x_0}{A}}{(1+\delta)x_0} - \eta(1+\delta)^{-2} \Big(1 + \frac{\log(A/x_0)}{\log(1+\delta)}\Big).$$

Proof of lemma C.2. For every $g = 0, 1, 2, \cdots$, let τ_g be the first iteration such that $x_t \ge (1+\delta)^g x_0$. Let b be the smallest integer such that $(1+\delta)^b x_0 \ge A$. By the definition of τ_g, we have $x_t \in [(1+\delta)^g x_0, (1+\delta)^{g+1} x_0)$ for all $t \in [\tau_g, \tau_{g+1})$ and $x_{\tau_{g+1}} \ge (1+\delta)^{g+1} x_0$, $x_{\tau_g - 1} < (1+\delta)^g x_0$, leading to

$$\begin{aligned} \sum_{t \in [\tau_g, \tau_{g+1})} \eta \cdot C_t [(1+\delta)^g x_0]^2 \le x_{\tau_{g+1}} - x_{\tau_g} &= \sum_{t \in [\tau_g, \tau_{g+1})} (x_{t+1} - x_t) \\ &= \sum_{t \in [\tau_g, \tau_{g+1})} \eta \cdot C_t x_t^2 \le \sum_{t \in [\tau_g, \tau_{g+1})} \eta \cdot C_t [(1+\delta)^{g+1} x_0]^2, \end{aligned}$$

following lower bound for $x_{\tau_{g+1}} - x_{\tau_g}$:

$$\begin{aligned} x_{\tau_{g+1}} - x_{\tau_g} = x_{\tau_{g+1}} - x_{\tau_g - 1} - \eta \cdot x_{\tau_g - 1}^2 &\ge (1+\delta)^{g+1} x_0 - (1+\delta)^g x_0 - \eta[(1+\delta)^g x_0]^2 \\ &= \delta(1+\delta)^g x_0 - \eta(1+\delta)^{2g} x_0^2, \end{aligned}$$

and following upper bound for $x_{\tau_{g+1}} - x_{\tau_g}$:

$$\begin{aligned} x_{\tau_{g+1}} - x_{\tau_g} = x_{\tau_{g+1} - 1} + \eta \cdot x_{\tau_{g+1} - 1}^2 - x_{\tau_g} &\le (1+\delta)^{g+1} x_0 + \eta[(1+\delta)^{(g+1)} x_0]^2 - (1+\delta)^g x_0 \\ &= \delta(1+\delta)^g x_0 + \eta(1+\delta)^{2(g+1)} x_0^2. \end{aligned}$$

Therefore,

$$\sum_{t \in [\tau_g, \tau_{g+1})} \eta \cdot C_t [(1+\delta)^g x_0]^2 \le \delta(1+\delta)^g x_0 + \eta(1+\delta)^{2(g+1)} x_0^2,$$

$$\sum_{t\in[\tau_g,\tau_{g+1})} \eta\cdot C_t[(1+\delta)^{g+1}x_0]^2 \geq \delta(1+\delta)^g x_0 - \eta(1+\delta)^{2g}x_0^2.$$

These imply that

$$\sum_{t\in[\tau_g,\tau_{g+1})} \eta\cdot C_t \leq \frac{\delta}{(1+\delta)^g x_0} + \eta(1+\delta)^2,$$

$$\sum_{t\in[\tau_g,\tau_{g+1})} \eta\cdot C_t \geq \frac{\delta}{(1+\delta)^{g+2} x_0} - \eta(1+\delta)^{-2}.$$

Recall b is the smallest integer such that $(1+\delta)^b x_0 \geq A$, so we can calculate that

$$\begin{aligned}\sum_{t\geq 0, x_t\leq A} \eta\cdot C_t &\leq \sum_{g=0}^{b-1} \frac{\delta}{(1+\delta)^g x_0} + \eta(1+\delta)^2 b = \frac{\delta\big(1-(1+\delta)^{-b}\big)}{\big(1-(1+\delta)^{-1}\big)x_0} + \eta(1+\delta)^{-2}b \\ &\leq \frac{\delta}{(1-(1+\delta)^{-1})x_0} + \eta(1+\delta)^2 b = \frac{1+\delta}{x_0} + \eta(1+\delta)^2 b\end{aligned}$$

and

$$\begin{aligned}\sum_{t\geq 0, x_t\leq A} \eta\cdot C_t &\geq \sum_{g=0}^{b-1} \frac{\delta}{(1+\delta)^{g+2} x_0} - \eta(1+\delta)^{-2} b = \frac{\delta\big(1-(1+\delta)^{-b}\big)}{(1+\delta)^2\big(1-(1+\delta)^{-1}\big)x_0} - \eta(1+\delta)^{-2}b \\ &\geq \frac{\delta\big(1-\frac{x_0}{A}\big)}{(1+\delta)^2\big(1-(1+\delta)^{-1}\big)x_0} - \eta(1+\delta)^{-2}b = \frac{1-\frac{x_0}{A}}{(1+\delta)x_0} - \eta(1+\delta)^{-2}b\end{aligned}$$

Note that $(1+\delta)^{b-1}x_0 < A$, i.e. $b \leq 1 + \frac{\log(A/x_0)}{\log(1+\delta)}$, therefore

$$\sum_{t\geq 0, x_t\leq A} \eta\cdot C_t \leq \frac{1+\delta}{x_0} + \eta(1+\delta)^2\Big(1+\frac{\log(A/x_0)}{\log(1+\delta)}\Big),$$

$$\sum_{t\geq 0, x_t\leq A} \eta\cdot C_t \geq \frac{1-\frac{x_0}{A}}{(1+\delta)x_0} - \eta(1+\delta)^{-2}\Big(1+\frac{\log(A/x_0)}{\log(1+\delta)}\Big).$$

□